arXiv:1407.0733v2 [cs.CV] 3 Oct 2014

Cortical spatio-temporal dimensionality reduction for visual grouping

Giacomo Cocci, Davide Barbieri, Giovanna Citti, Alessandro Sarti

October 6, 2014

Abstract

The visual systems of many mammals, including humans, is able to integrate the geometric information of visual stimuli and to perform cognitive tasks already at the first stages of the cortical processing. This is thought to be the result of a combination of mechanisms, which include feature extraction at single cell level and geometric processing by means of cells connectivity. We present a geometric model of such connectivities in the space of detected features associated to spatio-temporal visual stimuli, and show how they can be used to obtain low-level object segmentation. The main idea is that of defining a spectral clustering procedure with anisotropic affinities over datasets consisting of embeddings of the visual stimuli into higher dimensional spaces. Neural plausibility of the proposed arguments will be discussed.

Contents

1 Introduction **2**

2 The geometry of V1 **4**
2.1 The cortical feature space 4
2.2 Connectivity as a differential constraint 5
2.3 Discrete connectivity kernels 8

3 Spectral analysis of connectivities **10**
3.1 Spectral clustering 11
3.1.1 On possible neural implementations of the algorithm 14
3.2 Cortical affinity matrices 15
3.2.1 Reciprocal connectivities for spatially distributed features 15
3.2.2 Time asymmetry: a directed graphs approach 18

4 Visual grouping with cortical affinities **19**
4.1 Grouping with spatial features 20
4.2 Spatio-temporal grouping 24

5 Conclusions **29**

1 Introduction

It is well understood from the psychological theory of the Berliner Gestalt that local properties of the visual stimulus, such as neighboring, good continuation and common fate have a central role in the execution of global visual tasks like image segmentation and grouping [61, 57].

A key concept for the understanding of visual perceptual tasks is that of *association fields*, that was introduced in [17] to describe the sructure of the field of good continuation, underlying the recognition of perceptual units in the visual space. These results have been obtained by psychophysical experiments, presenting stimuli made of an ensemble of oriented patches, a subset of which was consistently aligned along a continuous path. The study of these phenomena allowed to identify the properties that a stimulus should have near a given patch in order to recognize such sample as belonging to a curvilinear path, namely co-linearity and co-circularity. The detection of these properties is indeed compatible with the functional behavior of simple cells in the primary visual cortex V1 as linear feature detectors for local orientations [23].

Several physiological experiments showed how the principles of association fields seem to be implemented in the V1 of mammals, where long-range horizontal connections preferentially link columns of neurons having similar preferred orientation [6]. On the other hand, by interpreting cortical columns as directional differential operators, in [9] it is shown how this specialized functional organization of V1 naturally leads to a geometric model of the association fields. The field lines are modeled with a family of integral curves on a contact structure based on the Lie algebra of the group of rigid motions of the Euclidean plane $SE(2)$. This geometric approach lies within a well established research line founded by the seminal works [27, 21, 35, 40], whose current state of the art can be found in [10] and references therein.

Further phenomenological experiments demonstrated the central role for the perception of global shapes of the features of movement direction and velocity [42] and that, similarly to what happens for the integration of spatial visual information, the brain is capable to predict complex stimulus trajectories, and to group together elements having similar motion or apparent motion paths [53, 54, 58]. Also in this case, the detection of such features is performed at the level of V1, where specialized cells show spatio-temporal behaviors optimized for the detection of local velocities [14, 11].

The analysis of spatio-temporal properties and organization of cortical visual neurons, together with the indications given by experimental results on visual spatial and motion integration, have recently led to extensions of the $SE(2)$ model which include local stimulus velocity. In [3], new classes of spatio-temporal connectivities were indeed introduced, providing a geometric model of association fields in the five dimensional contact structure of cells positions and activation times, together with the locally detected features of orientations and velocities. Such a structure embeds the purely spatial geometry in a layered fashion, and coherently integrates the association mechanisms as extensions to a higher dimensional space. Their capabilities in the elaboration of trajectories, and in the tasks of spatio-temporal image completion, actually showed good accordance with previous psychophysical and physiological results.

The aim of the present paper is to study the capabilities of such geometric spatio-temporal connectivities with respect to the tasks of image segmentation and grouping. Such tasks are addressed with spectral analysis, and in particular we will discuss and refine the proposed connectivity structures in order to properly describe them as oper-

ators on high dimensional feature spaces. The approach we will follow is to represent a spatio-temporal stimulus as a dataset in a feature space, and consider it as a weighted graph whose affinity matrix is determined by the geometric connectivity. The grouping mechanisms that we will describe arise from a spectral clustering of the associated graph Laplacian, making use of the probabilistic framework introduced in [33] followed by an adaptation of the simple and robust clustering technique proposed in [25] and, when dealing with nonsymmetric affinities, we followed the ideas introduced in [38]. We have tried to stuck with minimal hypotheses at the algorithmic level in order to keep focus on the role of the kernels. Our main results will show that the introduced spatio-temporal geometry provides connectivities suitable to robustly group spatio-temporal stimuli, but also that a connectivity pattern based also on local stimulus velocity can enhance the spatial grouping capabilities of a visual system.

With respect to actual neural implementations of such principles, apart from the study in [3], we can see that psychophysical experiments such as those of [18, 20, 30] addressed the problems of existence of association fields for local directions of motion and their role in visual grouping, and of comparison of human grouping with co-circular correlations in natural image statistics. Moreover, recent results of [48] show how spectral analysis of the spatial connectivity introduced in [9], that is the more basic one studied in the present paper, is actually implemented by the neural population dynamics of primary visual cortex. This suggests a fundamental role of spectral mechanisms in the phenomenology of perception, indicating that they may be concretely performed by the visual system, and hence providing a stronger motivation for the present detailed spectral analysis of connectivities and their segmentation properties of visual stimuli.

Finally, we would like to recall also that while co-circularity is naturally implemented in the presently studied kernels (see e.g. [47]), its use for the definition of affinity matrices whose spectra could be suited for line grouping was suggested already in [39].

The plan of the paper is the following. We will start in Section 2 with a description of the geometry arising from the spatio-temporal functional architecture of the visual cortex, as introduced in [2, 3], providing more detailed arguments on how to construct and implement the connectivity kernels we will use. In Section 3 we will recall the needed notions of spectral analysis on graphs, and introduce the basic clustering principles that we will adopt. Then we will show how to construct affinity matrices based on the spatio-temporal cortical geometry, providing specific discussion on how to deal with the different kinds of asymmetries present in the connectivities. Or main results will constitute Section 4, where we will use the previously introduced spectral clustering algorithm to automatically extract perceptual information from artificial stimuli living in the cortical feature spaces described in Section 2. In particular, we will propose two different connectivity mechanisms to perform spatio-temporal segmentation of motion contours and shapes, also providing parametric evaluations of the kernel performances, and we will discuss the relations with neural processes studied in several psychophysiological experiments.

Acknowledgements D. Barbieri was supported by a Marie Curie Intra European Fellowship (Prop. N. 626055) within the 7th European Community Framework Programme.

2 The geometry of V1

In this section we present a model of the functional architecture of the visual cortex as a contact structure, where the cortical connectivity is implemented as a diffusion process along its admissible directions. This same approach was taken in [3], where it was compared with psychophysical and physiological behaviors of the visual system.

2.1 The cortical feature space

It is well known since the fundamental studies of Hubel and Wiesel [23, 22] that primary visual cortex (V1) is one of the first physiological layers along the visual pathway to carry out geometrical measurements on the visual stimulus, decomposing it in a series of local feature components. The development of suitable electrophysiological techniques [43] has made possible to reconstruct the linear filtering behavior of V1 simple and complex cells, i.e. their spatio-temporal receptive profiles (RPs).

The RPs of orientation-selective cells in V1 have classically been modeled with two-dimensional Gabor functions [24], which basically compute a local approximation of the directional derivative of the visual stimulus, minimizing the uncertainty between localization in position and spatial frequency [13]. On the other hand, spatio-temporal RPs of V1 simple cells can be modeled by 3-dimensional Gabor functions of the form [11]

$$g^{\sigma}_{q,p}(x,y,t) = e^{2\pi i(p_x(x-q_x)+p_y(y-q_y)+p_t(t-q_t))} e^{-\left(\frac{(x-q_x)^2}{2\sigma_x^2}+\frac{(y-q_y)^2}{2\sigma_y^2}+\frac{(t-q_t)^2}{2\sigma_t^2}\right)}, \tag{1}$$

where $q=(q_x,q_y,q_t)$ is the spatio-temporal center of the Gabor filter, $p=(p_x,p_y,p_t)$ is the spatio-temporal frequency and $\sigma=(\sigma_x,\sigma_y,\sigma_t)$ is the spatio-temporal width. One of the crucial features of (1) is minimization of the uncertainty of simultaneous measurements in space-time and frequency. It is worth noting that this model strictly captures the features of so-called inseparable RPs, tuned for position, orientation and motion detection, depicted in Fig. 1 left. On the other hand, separable RPs can be obtained as linear superpositions.

Further analyses have also shown that the Gabor parameter distribution found in cortical cells cover only a subset of the whole Gabor family. Such subsets are optimized for the detection of the local features of orientation θ and speed v [11, 2]

$$\begin{aligned} \theta &= \arctan\left(\frac{p_y}{p_x}\right) \\ v &= \frac{p_t}{\sqrt{p_x^2+p_y^2}} \end{aligned}$$

that can be interpreted as fundamental features of the visual stimulus. For this reason we will not deal with the dependence on the spatial frequency $\kappa=\sqrt{p_x^2+p_y^2}$ or on the scales σ, but consider RPs of the form

$$g^{\sigma,\kappa}_{q,\theta}(x,y,t) = e^{-2\pi i\kappa((x-q_x)\cos\theta+(y-q_y)\sin\theta+v(t-q_t))} e^{-\left(\frac{(x-q_x)^2}{2\sigma_x^2}+\frac{(y-q_y)^2}{2\sigma_y^2}+\frac{(t-q_t)^2}{2\sigma_t^2}\right)} \tag{2}$$

where κ and σ are considered as fixed. This corresponds to a neural processing stage where the visual stimulus is lifted from the spatio-temporal image space $\mathbb{R}^2\times\mathbb{R}^+$ to the extended 5-dimensional feature space

$$\mathcal{M}_T = \mathbb{R}^2\times\mathbb{R}^+\times S^1\times\mathbb{R}^+ = \{\eta=(q_x,q_y,q_t,\theta,v)\},$$

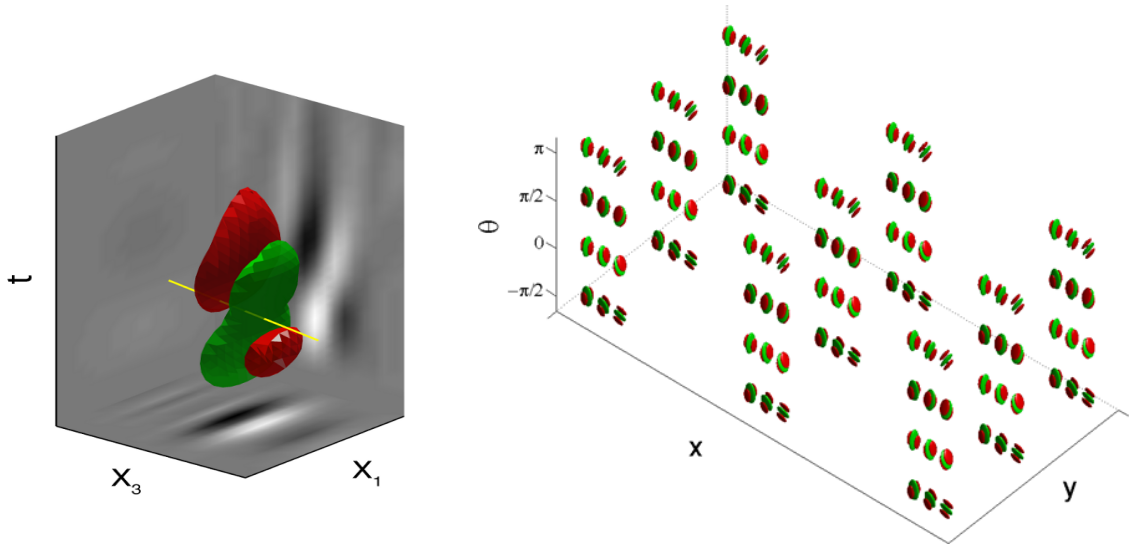

Figure 1: Left: isolevel surfaces of a V1 inseparable receptive profile depicted in space-time. Green and red surfaces enclose, respectively, excitatory and inhibitory regions. The yellow line indicates the direction of the local vector $\vec{X}_{\theta,v}$ (see text for details). Right: schematization of the feature-wise organization of the primary visual cortex. For each spatio-temporal point (x, y, t) of the image hyperplane there is a two dimensional fiber of representing local orientation θ and local velocity v.

where every point $\eta \in \mathcal{M}_T$ corresponds to a filter (2), as in Fig. 1 right. The activity of V1 simple cells is then considered as modeled by the map

$$f \mapsto F^{\sigma,\kappa}(q,\theta,v) \doteq \langle g^{\sigma,\kappa}_{q,\theta}, f\rangle_{L^2(\mathbb{R}^3)}. \tag{3}$$

2.2 Connectivity as a differential constraint

The functional behavior of V1 simple cells modeled by (3) can be interpreted as a finite scale spatio-temporal directional derivative of the stimulus around position q, performed along the direction

$$\vec{X}_{\theta,v} = (\cos\theta, \sin\theta, -v)$$

expressed in the coordinates $\{\hat{e}_x, \hat{e}_y, \hat{e}_t\}$. Accordingly, this derivation is maximal along the direction of the gradient of f. This implies that the lifting to $\mathcal{M}_T$ of any smooth level set of f is always orthogonal to the vector field

$$\vec{X} = (\cos\theta, \sin\theta, -v, 0, 0) \in \mathcal{T}_\eta \mathcal{M}_T$$

expressed in the coordinates $\{\hat{e}_{q_x}, \hat{e}_{q_y}, \hat{e}_{q_t}, \hat{e}_\theta, \hat{e}_v\}$, where $\mathcal{T}_\eta\mathcal{M}_T$ stands for the tangent space of $\mathcal{M}_T$ at $\eta = (q_x, q_y, q_t, \theta, v)$.

Hence the present problem induces to consider as admissible surfaces on $\mathcal{M}_T$ those whose tangent space at any point is spanned by the vector fields

$$\begin{aligned} \vec{X}_1 &= (-\sin\theta, \cos\theta, 0, 0, 0)\,, \;\; \vec{X}_2 = (0,0,0,1,0) \\ \vec{X}_4 &= (0,0,0,0,1)\,, \;\; \vec{X}_5 = (v\cos\theta, v\sin\theta, 1, 0, 0) \end{aligned} \tag{4}$$

defining the orthogonal complement to $\vec{X}$ in $\mathcal{T}\mathcal{M}_T$.

The four dimensional hyperplanes generated by $\{X_1, X_2, X_4, X_5\}$ is called contact plane, and the whole structure is named a contact structure. Contact structures have been used for modeling the functional architecture of the visual cortex in several works, see e.g. [40, 9].

Due to this contact structure, the connectivity among V1 cells on $\mathcal{M}_T$ can be modeled geometrically as in [3] in terms of advection-diffusion processes along the directions of the vector fields (4). Two corresponding stochastic processes were introduced in order to provide concrete realizations of the mechanisms of propagation of information along connections.

A first mechanism, aimed to model connectivity along lifted contours of a spatial image at a fixed time, consists of a propagation along the direction $\vec{X}_1$ forced by a diffusion over $\vec{X}_2$ and $\vec{X}_4$. This will be used for a single-frame segmentation out of a spatio-temporal streaming. It lives on a codimension 1 submanifold of $\mathcal{M}_T$ at fixed time t_0 that we will call $\mathcal{M}_0 = \mathbb{R}^2 \times S^1 \times \mathbb{R}^+ = \{\xi = (q_x, q_y, q_t, \theta, v) : q_t = t_0\}$, and can be formally described by the following system of stochastic differential equations

$$\begin{cases} dq_x(s) &= -\sin\theta(s)\, ds \\ dq_y(s) &= \cos\theta(s)\, ds \\ d\theta(s) &= \kappa\ dW_1(s) \\ dv(s) &= \alpha\ dW_2(s) \\ \xi(s=0) &= \xi_0 \end{cases} \tag{5}$$

where $W = (W_1, W_2)$ is a two dimensional Brownian motion and (κ, α) are the corresponding diffusion constants. The Fokker-Plank equation associated to this process, which provides a transition probability density ρ_0, is

$$(\partial_s - L_0)\rho_0(\xi, s|\xi_0, 0) = \delta(\xi - \xi_0)\delta(s) \tag{6}$$

where the evolution operator is given by $L_0 = X_1 - \kappa^2 X_2^2 - \alpha^2 X_4^2$ and, as customary, we have denoted with X_i the directional derivative along the vector field $\vec{X}_i$.

A second mechanism, aimed to model connectivity among moving contours of a spatio-temporal stimulus, will be used for spatio-temporal segmentation of apparent point trajectories. It consists of a propagation along $\vec{X}_5$ again forced by a diffusion over $\vec{X}_2$ and $\vec{X}_4$, and is described by the stochastic process on $\mathcal{M}_T$

$$\begin{cases} dq_x(s) &= v\cos\theta(s)\, ds \\ dq_y(s) &= v\sin\theta(s)\, ds \\ dq_t(s) &= ds \\ d\theta(s) &= \kappa\ dW_1(s) \\ dv(s) &= \alpha\ dW_2(s) \\ \eta(s=0) &= \eta_0. \end{cases} \tag{7}$$

The Fokker-Plank equation associated to this process, which provides a transition probability density ρ_T, is

$$(\partial_s - L)\rho_T(\eta, s|\eta_0, 0) = \delta(\eta - \eta_0)\delta(s) \tag{8}$$

where the evolution operator is given by $L = X_5 - \kappa^2 X_2^2 - \alpha^2 X_4^2$.

The structure of these processes explicitly assigns different roles to the spatio-temporal variables q, where the stimulus is defined, and to the engrafted variables (θ, v). More precisely, we have advection in the q variables while diffusion occurs in the (θ, v) variables. It is worth noting that this construction naturally extend the process proposed by Mumford in [35] for the case of static images, which consists of a propagation along the direction $\vec{X}_1$ forced by a diffusion over $\vec{X}_2$, on the three dimensional submanifold $\mathcal{M}_3 = \mathbb{R}^2 \times S^1 = \{\zeta = (q_x, q_y, \theta)\}$

$$\begin{cases} dq_x(s) &= -\sin\theta(s)\, ds \\ dq_y(s) &= \cos\theta(s)\, ds \\ d\theta(s) &= \kappa\, dW_1(s) \\ \zeta(s=0) &= \xi_0. \end{cases} \tag{9}$$

The Fokker-Plank equation associated to this process, which provides a transition probability density ρ_3, is

$$(\partial_s - L_3)\rho_3(\zeta, s|\zeta_0, 0) = \delta(\zeta - \zeta_0)\delta(s) \tag{10}$$

where the evolution operator is given by $L_3 = X_1 - \kappa^2 X_2^2$.

Each of the three equations (6), (8) and (10) is defined by a Markov generator $\mathcal{L}$ of a stochastic process over a manifold X, with transition probability ϱ satisfying

$$(\partial_s - \mathcal{L})\varrho(x, s|x_0, 0) = \delta(x - x_0)\delta(s), \tag{11}$$

and our aim is to use the density ϱ to define a connectivity kernel over the manifold X, disregarding the dynamics over the evolution parameter s. This connectivity be obtained by integrating ϱ over s, against some appropriately chosen weight p, hence defining a connectivity kernel as

$$\Gamma(x, x_0) = \int_0^\infty \varrho(x, s|x_0, 0)p(s)ds.$$

When p is a positive weight normalized to 1, this kernel has the well known probabilistic interpretation of being the transition probability for a stochastic process subordinated to a time process with independent increments distributed with p (see e.g. [49]). One common choice for the weight p is that of [35, 62], that is an exponential decay representing a decay of the signal during the propagation, which amounts to replace ϱ with its Laplace transform at a fixed time-scale. Another possible choice is the one used in [47, 3] where p was identically set to 1, which amounts to consider the density of points reached at any value of the evolution parameter. In this work we will deal with an intermediate choice between these two, namely we will choose a weight depending on an evolution-scale parameter H as

$$p(s) = \frac{1}{H}\chi_{[0,H]}(s) = \begin{cases} \frac{1}{H} & s \in [0, H] \\ 0 & \text{otherwise.} \end{cases}$$

This choice amounts to assign a uniform weight to all values of the evolution parameter, but allows to keep track of the evolution length over which the stochastic paths are evaluated. Our notion of connectivity kernel will then be

$$\Gamma^H(x, x_0) = \frac{1}{H}\int_0^H \varrho(x, s|x_0, 0)ds. \tag{12}$$

It is worth noting that both the diffusion parameter κ and the evolution-scale parameter H could be treated, in principle, as additional fiber variables of the model, linking their different associated connectivities to some prominent feature of the image detected by the visual cortex, for example curvature and scale. For the sake of model simplicity, though, in this paper we will treat these variables as parameters, addressing the extension of the model in a future work.

2.3 Discrete connectivity kernels

In this section we outline the numerical method we have used to compute the connectivity kernels Γ^H, which are needed due to the lack of analytic solutions. A notable exception is the $SE(2)$ case (10), for which an analytic solution was obtained in [15] and compared in [65] to several numerical approaches, all of them being different from the present one. Our choice of numerics has privileged the robustness and flexibility of the implementation, which has been applied to all the three cases.

The differential equations of type (11) that we need to solve originate from a stochastic process over X, being this $\mathcal{M}_0$, $\mathcal{M}_T$ or $\mathcal{M}_3$, that we can write in terms of its sample paths $\gamma : \mathbb{R}^+ \to X$ as

$$d\gamma(s) = A(\gamma)ds + B(\gamma)dW$$

where A is a vector field and B is a matrix field over X representing, respectively, (5), (7) and (9). A flexible and efficient numerical technique is then that of Markov Chain Monte Carlo methods (MCMC) (see e.g. [19]), which was implemented as follows. We first fix a discretization over the parameter s, using without loss of generality a step $\Delta s = 1$ so that the discrete evolution will be performed over $\mathbb{N}$, and a discrete covering grid $\{\Omega_j\}_{j\in\mathbb{N}}$ of X, i.e. a collection of subsets of X satisfying $\Omega_i \cap \Omega_j = \emptyset$ if $i \neq j$ and $\bigcup \Omega_j = X$. For a given $\gamma_0 \in X$ we then simulate N several discrete-time random paths over X with the recursive equation

$$\gamma_{h+1} = \gamma_h + A(\gamma_h) + B(\gamma_h)\delta_h, \quad h \in \mathbb{N}$$

where $\{\delta_h\}_{h\in\mathbb{N}}$ are (vector valued) i.i.d. gaussian random variables, and assign to each region Ω_j a value between 0 and 1 corresponding to the number of paths that passed through it divided by N. This provides a distribution over the cells Ω_j that, up to a multiplicative constant, for large values of N gives a discrete approximation of the solution to (11) that we will denote $\varrho(\Omega_j, h|\gamma_0, 0)$. The resulting discrete approximation of the connectivity kernel (12) will then be computed as

$$\hat{\Gamma}^H(\Omega_j|\gamma_0) = \frac{1}{H}\sum_{h=0}^{H} \varrho(\Omega_j, h|\gamma_0, 0). \tag{13}$$

A deeper discussion of this kind of numerical approximation for stochastic differential equations can be found in [41] and references therein.

During the simulations of Section 4 we will vary the parameters which characterize the kernels. In order to keep track of their dependency we will then use the notation $\hat{\Gamma}_0^{H,\kappa,\alpha}$ and $\hat{\Gamma}_T^{H,\kappa,\alpha}$ for the discrete approximations of the connectivity kernels in the cases, respectively, of (5) and (7), where κ and α stand for the diffusion coefficients over θ and v, and the notation $\hat{\Gamma}_3^{H,\kappa}$ for the discrete approximation of the connectivity kernel in the case (9), where κ is the diffusion coefficient over θ.

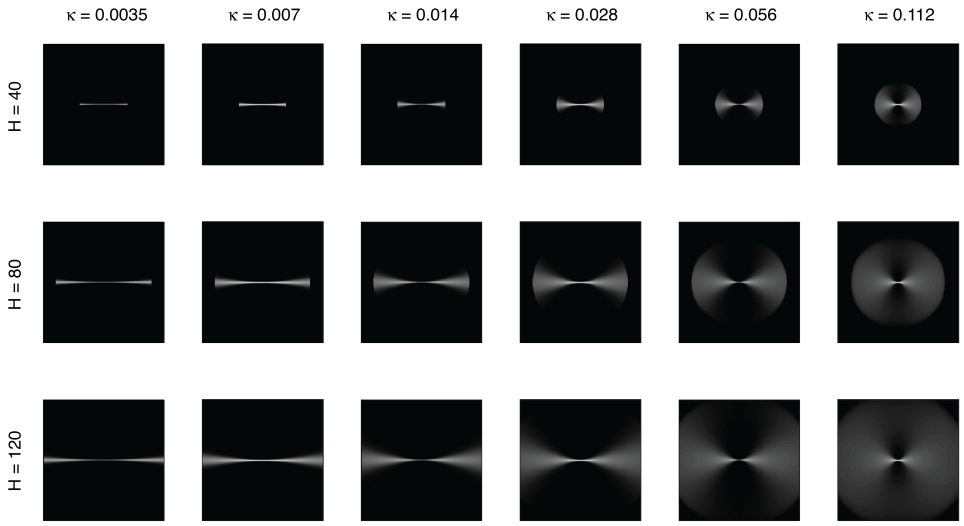

Figure 2: Marginal distributions over the (q_x, q_y) plane of the kernel $\hat{\Gamma}_3^{H,\kappa}$ computed for different values of the parameters k and H.

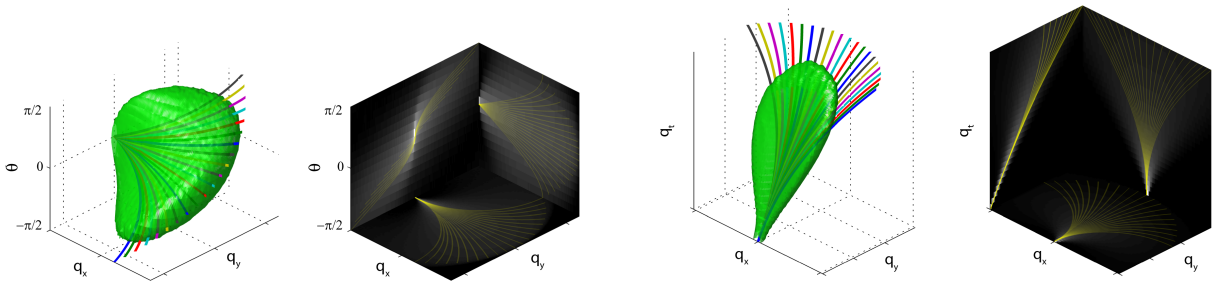

Figure 3: Stochastic kernels $\hat{\Gamma}_0^{H,\kappa,\alpha}$ (top) and $\hat{\Gamma}_T^{H,\kappa,\alpha}$ (bottom) for $N = 10^6$ stochastic paths, compared to the horizontal curves calculated as in [3]. Left: isosurface plot of the kernels. Right: kernel projections relative to the variables (q_x, q_y, θ) for $\hat{\Gamma}_0^{H,\kappa,\alpha}$ and (q_x, q_y, q_t) for $\hat{\Gamma}_T^{H,\kappa,\alpha}$ (gray), under the projections of the horizontal curves (yellow).

The effect of the variation of the parameter κ on the shape of $\hat{\Gamma}_3^{H,\kappa}$ is shown in Fig. 2, while in Fig. 3 we show the isosurfaces of two projections of the connectivity kernels $\hat{\Gamma}_0^{H,\kappa,\alpha}$ and $\hat{\Gamma}_T^{H,\kappa,\alpha}$ together with the related *horizontal curves*, i.e. the curves obtained from the systems (5) and (7) by substituting noise with a constant, and varying the parameters κ and α. These curves were introduced in [3] as the geometric counterpart, in the space-time contact structure, to the horizontal curves used in [9] to model the orientation association fields. They provide then a natural geometric extension of the notion of association fields to spatio-temporal stimuli with orientation and velocity features. As it can be seen, the related kernels reach their maximum values in the proximity of the fan of such curves originating from the same starting point, which motivates their role as a model for a spatio-temporal neural connectivity.

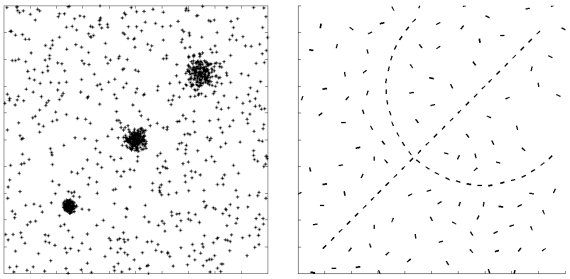

Figure 4: Sample spatial stimuli containing perceptual units embedded in random fields.

3 Spectral analysis of connectivities

In this section we recall the main notions to be used about clustering with spectral analysis on graphs, and how to use the previously introduced geometric setting for these purposes. The task we will address fall into the well-known class of problems broadly referred to as dimensionality reduction, which deals with the general problems of data partitioning and of locality-preserving embeddings of high-dimensional data sets. The literature on these topics is huge, and we have chosen to refer only to works directly related to the present one, addressing the reader to the bibliography contained therein.

The cognitive task of spatial or spatio-temporal visual grouping can be interpreted as a form of clustering. Two examples of spatial visual stimuli providing standard clustering problems are portrayed in Fig. 4. On the left, three dense Gaussian distributions of 2-dimensional points are embedded within a sparser set of random points uniformly scattered throughout the domain. The human visual system normally segregates the points of the Gaussian clouds into three separate groups of points (objects) lying on a noisy environment. On the right, two dashed continuous lines are embedded in a field of segments having random position and orientation. In this case, stimulus collinearity gives rise to a pop-out effect that makes the two lines easily distinguishable from the background, that is the phenomenon quantified by the psychophysical experiments of [17].

Although these examples show two quite different grouping effects, a common underlying mechanism can be formalized as follows. Given a data set of n points $S = \{x_i\}_{i=1}^n \subset X$ living in an arbitrary metric space (X, d), the task of grouping together the points that – according to the distance d – are closer, or more *similar* to each other (so that their ensemble forms an *object*), amounts to identify K disjoint subsets $S_i \subset S$, with $i \in \{1, \dots, K\}$, where $K - 1$ of them contain points relatively close to each other and relatively far from the rest, and one of them contains points that cannot be classified in such a way, and that will be considered as noise.

It is widely known that this problem is not easily resolvable with purely clustering algorithms like *K-means* (see e.g. [36]). A major branch of research has been recently devoted to the development of spectral techniques that allow to address these issues in terms of the spectral properties of symmetric positive semi-definite affinity matrices constructed from the input data set. The ensemble of these techniques can generally be subdivided into two classes [28]: methods for locality-preserving embeddings of large data sets, that project the data points onto the eigenspaces of the affinity matrices [46, 5, 12], and methods for data segregation and partitioning, that perform an additional clustering step taking as input the projected data set (see e.g. [39, 59, 51, 33]).

In Subsection 3.1 we will describe the spectral clustering algorithm we will use to perform visual grouping, while in Subsection 3.2 we will show how to use the geometric feature spaces and the cortical connectivities developed in Section 2 for the construction of affinity matrices associated to spatial and spatio-temporal visual stimuli. This method will be applied in Section 4 to several stimuli.

3.1 Spectral clustering

Let us consider the data set $S = \{x_i\}_{i=1}^n$ in the metric space (X, d) as the vertices of a weighted graph, where the edge weights $\{a_{ij}\}_{i,1=1}^n$ define an *affinity matrix* A. It has been originally shown in [39] that, when A is a real symmetric matrix, its first eigenvector of can serve as indicator vector for basic grouping purposes. In the same work, the authors also proposed a partitioning algorithm that recursively separate the foreground information from the data set. While this algorithm's implementation is straightforward and efficient, it was demontrated that it can easily lead to clustering errors due to noise, non-linear distributions or outliers [59]. This argument was later improved, also in view of its relation to several other problems such as that of minimal graph cuts [51], essentially by performing the spectral analysis of a suitably normalized affinity matrix (see e.g. [56] and references therein). We will use the normalization proposed in [33], which turns a real symmetric affinity matrix A into the transition matrix P of a reversible Markov chain via row-wise normalization. More precisely, if D is the diagonal degree matrix, having elements

$$d_i = \sum_{j=1}^{n} a_{ij},$$

the normalized affinity matrix P is given by

$$P = D^{-1}A. \tag{14}$$

This matrix won't be, in general, symmetric, but it can be shown that its eigenvalues $\{\lambda_j\}_{j=1}^n$ are real and satisfy $0 \leq \lambda_j \leq 1$, and its eigenvectors $\{u_j\}_{j=1}^n$ can accordingly be chosen with real components. The clustering properties of the eigenvectors of P can be clearly understood in the following ideal case. Suppose that the graph $G = (S, A)$ with nodes given by S and edge weights given by A has K connected components $\{G_i\}_{i=1}^K$, and that all the elements of each component have the same edge weight connecting them. The resulting normalized affinity matrix P would then be a block diagonal matrix, with only K non-null eigenvalues $\{\lambda_i\}_{i=1}^K$ each of them equal to 1, whose corresponding eigenvectors $\{u_i\}_{i=1}^K$ are piece-wise constant indicator functions of the partitions.

In real applications, the affinity matrices are perturbed versions of the block diagonal ones and do not possess an ideal binary spectra with purely indicator eigenvalues, thus making the partitioning problem generally ill-posed. The best situation one can hope is then that of a good approximation of the ideal case, where the affinity on the dataset is such that there are clusters of points that are strongly connected mainly to their cluster neighbours, and only weakly connected to the rest. In general, several authors demonstrated that the egenvectors $\{u_i\}_{i=1}^K$ of the normalized affinity matrix P, corresponding to the K largest eigenvalues $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_K$, solve the relaxed optimization problem of normalized graph cuts, and give a nice probabilistic interpretation of the clustering problem.

A crucial step in real applications is then that of choosing the value of K, that is deciding how many eigenvalues are worth to be taken into consideration and consequently how many eigenvectors possess relevant clustering information. Many authors proposed different solutions, like looking for the maximum eigengap or trying to minimize a particular cost function (see e.g. [64]). We decided to adopt a semi-supervised solution consisting of fixing an a-priori significance threshold ϵ, and consider as clustering eigenvectors all those u_i whose $\lambda_i > 1 - \epsilon$.

Since the more P is far from being similar to a block diagonal matrix, the more its spectrum will be far from being dichotomous, with the ordered λ's decreasing more smoothly (see e.g. Fig. 5), in the more ill-posed cases the sensitivity to small changes on the values of ϵ may become very high. In order to facilitate this delicate passage, we have then used a technique suggested by the well-known diffusion map approach [12, 28]: we have introduced an auxiliary thresholding integer parameter τ, and we have evaluated the exponentiated spectrum $\{\lambda_i^\tau\}_{i=1}^n$, which for sufficiently large values of τ is closer to dichotomy, against the threshold ϵ. This exponentiated spectrum has an easy probabilistic interpretation, being the spectrum of the matrix P^τ. Indeed, due to the normalization (14), P can be seen as the Markov transition matrix of a random walk over the graph, so that P^τ represents the transition probability of the same random walk in τ steps.

Once the number K of eigenvectors to use has been selected, we have used a variation of a simple clustering technique proposed in [25] in order to extract the clustering information. It corresponds, in the notation of [25], to the clustering of the rows of P based on its reduced matrix of eigenvectors $\hat{U} = [u_1 \dots u_q]$. The main differences are that, once the thresholding parameters are fixed, this algorithm dynamically assigns the number of (pre)clusters, and that a notion of background is explicitly introduced in the last step. We have summarized our spectral clustering algorithm in Table 1.

1.	Build the affinity matrix A upon an appropriate connectivity measure
2.	Compute the normalized affinity matrix $P = D^{-1}A$
3.	Solve the eigenvalue problem $Pu_i = \lambda_i u_i$, where the order of indices $i = 1 \dots n$ is such that $\{\lambda_i\}_{k=1}^n$ is decreasing
4.	Fix a threshold ϵ and a diffusion parameter τ
5.	Define $q = q(P; \epsilon, \tau)$ as the largest integer such that $\lambda_i^\tau > 1 - \epsilon$ for all $i = 1 \dots q$
6.	Define the preclusters $\{\hat{C}_k\}_{k=1}^q \subset \mathbb{R}^d$ so that for any $i = 1 \dots n$ the point x_i belongs to $\hat{C}_k$ whenever $$k = \arg\max_{j \in \{1 \dots q\}} \{u_j(i)\}$$
7.	Fix a minimum cluster size M, join together the preclusters with less than M elements into the cluster C_0, and order the remaining preclusters so to obtain a partition of the dataset into $\{C_k\}_{k=1}^K \subset \mathbb{R}^d$, where $K = K(P; \epsilon, \tau, M)$.

Table 1: Spectral clustering algorithm.

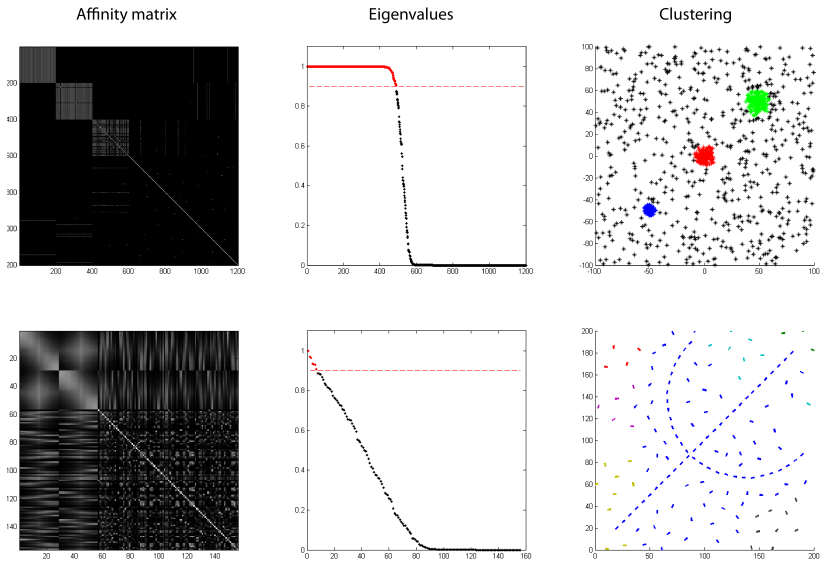

Figure 5: Result of the proposed algorithm for the two example data sets. The picture shows how an Euclidean isotropic kernel happens to be an optimal choice to cluster groups of points living in the image plane $\mathbb{R}^2$, but results inadequate when trying to separate boundaries or contours, whose additional feature of local orientation suggests to consider as points on the contact manifold $\mathcal{M}_3 = \mathbb{R}^2 \times S^1$.

As a first application of this algorithm, consider the two data sets presented in Fig. 4 as sets of points in $\mathbb{R}^2$ endowed with the isotropic Euclidean distance $d(x, y) = \|x-y\|$, and define the affinity matrix

$$a_{ij} = e^{-\frac{d(x_i,x_j)^2}{2\sigma^2}}, \tag{15}$$

where σ is a scale parameter that has to be chosen upon the characteristics of the data set to be clustered. The result is that of two connected graphs where the similarity between vertices decays as a Gaussian with their Euclidean distance. This is intuitively suitable to describe the visual clustering of the first dataset, but does not take into consideration the kind of similarity that characterizes the visual grouping of the second one. The results of the application of the spectral clustering algorithm with affinity given by (15) is shown in Fig. 5. In the first case, the algorithm performs the clustering process correctly, finding automatically the number of the main perceptual units and assigning the remaining elements to the noise/background cluster. It is worth noting that the spectrum of the normalized affinity matrix P^τ counts many eigenvalues that are close to 1, each only representing a single perceptual unit, that are mostly composed of few elements. The segment data set, on the other hand, was not clustered correctly: this was predictable, as these kind of stimuli are characterized by a local feature of orientation and hence should be better considered on the position-orientation domain $\mathbb{R}^2 \times S^1$, together with an anisotropic affinity.

3.1.1 On possible neural implementations of the algorithm

Besides the purely computational presentation that we have given of our spectral clustering algorithm, it may actually be relevant to consider whether plausible neural computations exist that could be responsible of its implementation in the visual cortex.

The discussion of Section 2 and the work [3] on the psychophysical and physiological nature of the connectivity kernels constitute our main motivation to consider the first step of the algorithm as a cortical implementation of affinities in the feature spaces. On the other hand, the normalization step 2 can be understood as formally analogous to the one introduced in [52], and may be motivated by the evidence that biological neural systems tend to adjust the weight of afferent connections so that a neuron with few incoming connections will weight those inputs more heavily than a neuron with many incoming connections (see also the discussion in [4]).

With respect to step 3, previous works such as [7, 16] describe a possible implementation of the spectral analysis as a mean-field neural computation, obtaining the emergence of eigenvectors of the connectivity as a symmetry breaking in the evolution equation associated to sufficiently high eigenvalues. An extension of these works can be found in [48], where it is shown that in the presence of a visual input the emerging eigenvectors correspond to visual perceptual units, which are thus obtained from a spectral clustering on excited connectivity kernels. According to these models, when the magnitude of an eigenvalue trespasses a given threshold, its corresponding eigenvector becomes a locally unstable solution, hence generating a distinguishable activity pattern. This threshold is what we aim to reproduce with steps 4 and 5.

Coming to step 6, we first note that its combinatorial formulation represents a concrete implementation of the following principle: eigenvectors with high eigenvalues represent preceptual units, and a point of the visual stimulus is assigned the unit whose eigenvector component over such a point has the higher magnitude. As we will discuss more in detail below, such a clustering will perform better the more the affinity matrix is close to a block diagonal matrix, since in such cases the eigenvectors relative to high eigenvalues will be close to indicator functions of weakly overlapping regions of the stimulus. In these situations the concrete neural implementation of steps 3, 4 and 5 already includes our step 6. Indeed, the computation of an almost-indicator eigenvector that is an unstable solutions to the neural field equation means that the underlying cells have been selected and participate to the associated activity pattern. We also observe at this point that in this normalized model all sufficiently high eigenvalues have very close magnitudes, that in the ideal case is 1 for all of them. If a neural field model with two populations is considered, the magnitude of the eigenvalue is associated to the frequency of oscillation of the neural population, which in this situations is almost the same for all perceptual unit and lies in the gamma range. This aspect marks a substantial difference with other non-normalized models of spectral neural computations.

The aim of step 7 is to measure the *saliency* of the perceptual units, and to introduce a thresholding on it. The notion of saliency that we consider is proportional to the total neural activity involved in the detection of a perceptual unit that, in the special artificial cases under consideration, corresponds to the number of related elements of the dataset.

3.2 Cortical affinity matrices

As we have described in Section 2, a spatial or a spatio-temporal visual stimulus is represented by different classes of V1 specialized neurons on different kinds of higher dimensional feature spaces, and we have focused on three of them that we have called $\mathcal{M}_0$, $\mathcal{M}_T$ and $\mathcal{M}_3$. The neural response (3) is in general non null on the whole feature space, but according to the stronger or weaker presence of the locally preferred feature, the magnitude of the response will be higher or lower. Here we will adopt a great simplification, and consider the response magnitude to be allowed to take only value 0 or 1, so that the datasets we will deal with will be purely point datasets embedded in a feature space. The main advantage of this simplification is the possibility to test the grouping capabilities of the geometric connectivities in the feature spaces without dealing with the many delicate questions related to the stimulus representation itself. On the other hand, only a few classes of stimuli are feasible to be represented in such a way. For this reason we will work only with synthetic stimuli and generate the corresponding datasets directly in the feature space. This indeed seems to be the safest way to separate the problems related to the geometric properties of the proposed connectivity from the ones related to the good representation of stimuli onto the feature space.

Let then X be a feature space, such as $\mathcal{M}_0$, $\mathcal{M}_T$ or $\mathcal{M}_3$, let $\{\Omega_j\}_{j\in\mathbb{N}}$ be a discrete covering such as the one introduced in Subsection 2.3, and let $S = \{x_i\}_{i=1}^n \subset X$ be the dataset representing the visual stimulus on this higher dimensional space, being this of type $S_0 = \{\xi_i\}_{i=1}^n \subset \mathcal{M}_0$, $S_T = \{\xi_i\}_{i=1}^n \subset \mathcal{M}_T$ or $S_3 = \{\zeta_i\}_{i=1}^n \subset \mathcal{M}_3$. Our aim is to define an affinity that makes two points x_i and x_j more similar the higher is their geometric connectivity, computed as in (12) with $\hat{\Gamma}^H(x_i, x_j)$. A necessary compatibility condition between the dataset and the discretization of the space is that each Ω_j contains at most one point x_i of the dataset. This indeed makes the kernels well defined kernels on the dataset, and allows us write $\hat{\Gamma}^H(x_i, x_j)$. However, since the associated Markov generator $\mathcal{L}$ is not selfadjoint, this kernel is not symmetric. We will then need to construct symmetric affinities from these connectivities, or rather adapt the previously described theoretical setting to couple with this asymmetry: we will choose one or the other way depending on the kernel, depending on the geometry of the information carried.

3.2.1 Reciprocal connectivities for spatially distributed features

A reasonable neural assumption when modeling long-range horizontal connections among V1 cells on $\mathcal{M}_3$ or on $\mathcal{M}_0$, which refer to the spatially distributed features of local orientations and local velocities, is that these connections are reciprocal [6, 26]. This means that two cells that are selective for the features (θ_1, v_1) and (θ_2, v_2) are assumed to be symmetrically connected. We will then take as affinities $A^H = (a_{ij}^H)$ a symmetrization of the corresponding connectivity kernels. The symmetrization we have chosen consists of taking the Hermitian component of the kernels:

$$a_{ij}^H = \frac{\hat{\Gamma}^H(x_i, x_j) + \hat{\Gamma}^H(x_j, x_i)}{2}. \tag{16}$$

This symmetrization has indeed a specific geometric meaning, which can be easily understood for the kernel $\hat{\Gamma}_3^{H,\kappa}$: it is equivalent to sum the fundamental solution of (10) with the fundamental solution of the same operator under an angular shift of 180

degrees in the variable θ. Such rotation turns the drift term X_1 into $-X_1$, hence transforming the forward Kolmogorov equation into the corresponding backward equation, so the sum of the two solutions is clearly symmetric. On the other hand, that sum allows to identify angles up to 180 degrees, hence turning a process which was a priori defined over the group $SE(2)$ of positions and angles into a process properly defined on positions and orientations. For this reason, such symmetrization is customary when dealing with such nuclei with the purpose of describing pure orientation (see e.g. [47]). We observe that this symmetrization does not modify the single cell response of modeled simple cells, which still detect angles, but rather introduces a symmetry in their geometric connectivities so that a cell having a preferred angle is considered to be long range connected to other cells along the two angles corresponding to the associated orientation. The same argument applies to the kernel $\hat{\Gamma}_0^{H,\kappa,\alpha}$, whose Fokker-Planck operator differs only for an additional diffusion term on the velocities.

Depending on the application to clustering with the kernels $\hat{\Gamma}_0^{H,\kappa,\alpha}$ or $\hat{\Gamma}_3^{H,\kappa}$ we will then obtain from (16) different affinity matrices over datasets $\{\xi_j\} \subset \mathcal{M}_0$ or $\{\zeta_i\} \subset \mathcal{M}_3$. The associated spectral clustering will then depend on the parameters (ϵ, τ) and M from the algorithm, and on the parametes H, κ and α from the kernels.

The grouping capabilities of such anisotropic affinities, and the role of the parameters, may be first evaluated on the simplest connectivity $\hat{\Gamma}_3^{H,\kappa}$. In order to do so, we have applied the spectral clustering algorithm with the cortical affinity to the second stimulus of Fig. 4. Such stimulus was indeed considered as a dataset in $\mathcal{M}_3 = \mathbb{R}^2 \times S^1$, each segment having a position in $\mathbb{R}^2$ determined by its center, and a position in S^1 corresponding to its orientation. The results obtained are displayed in Fig. 6 for different sets of kernel parameters H and κ. The clustering parameters used here and in the examples of the next sections were $\epsilon = 0.05$, $\tau = 150$ and $M = 3$: we have kept those parameters fixed for all the tests, in order to focus on the differences in the grouping properties due to the connectivity kernels.

As we can see, the presence of an affinity that is more semantically and geometrically adapted to the dataset with respect to (15) positively influences the grouping capabilities of the method. However, already this first simulation shows the relevance of the kernel parameters on the quality of the grouping.

For the top plots of Fig. 6, we have used an evolution integration step of $H = 40$ and an orientation diffusion coefficient $\kappa = 0.014$, which coincides with the curvature of the semi-circular object. The algorithm clearly succeeds in distinguishing the two perceptual units from each other, and correctly assigns the remaining elements to the same background/noise partition. For the middle plots, we have reduced the value of the diffusion constant to $\kappa = 0.0035$: while the algorithm correctly retrieves the straight contour and distinguish the units from the background, the semi-circle gets over-partitioned. As in the previous case, the affinity matrix is close to being a block diagonal matrix, but in this case the partitioned connected components of the sub-graph containing the two objects are more than two. Setting again the diffusion coefficient to the original value, for the bottom plot we have increased the evolution integration step value to $H = 100$. As the stimulus domain is 200×200 pixel wide, this means that every segment in the example could potentially have a non-null connectivity value with almost an half of the other segments, if the co-circularity conditions are satisfied. Indeed, by observing the resulting affinity matrix $A_3^{H=100,\kappa=0.014}$, we can see high affinity values between the objects and the random noise and in between the background elements. Again, even if the straight line is correctly retrieved as a single object, two random

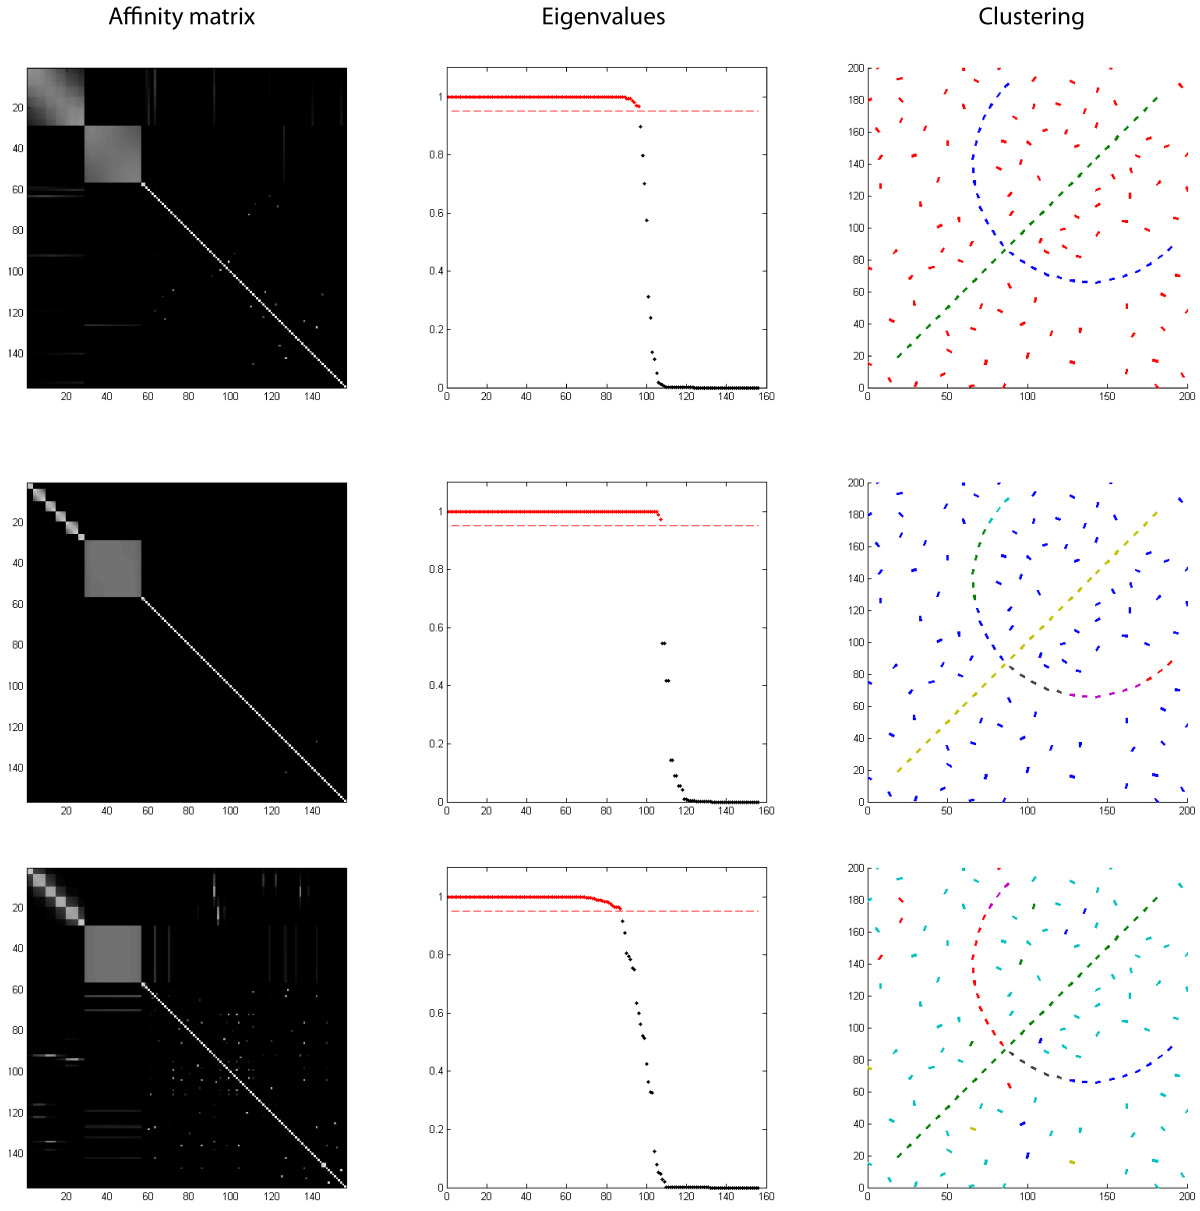

Figure 6: Result of the proposed algorithm for the second example data set and different parameters for the kernel $\hat{\Gamma}_3^{H,\kappa}$. Different kernel parameters modify the look of the affinity matrix, of its spectrum, and of the resulting data set partitioning.

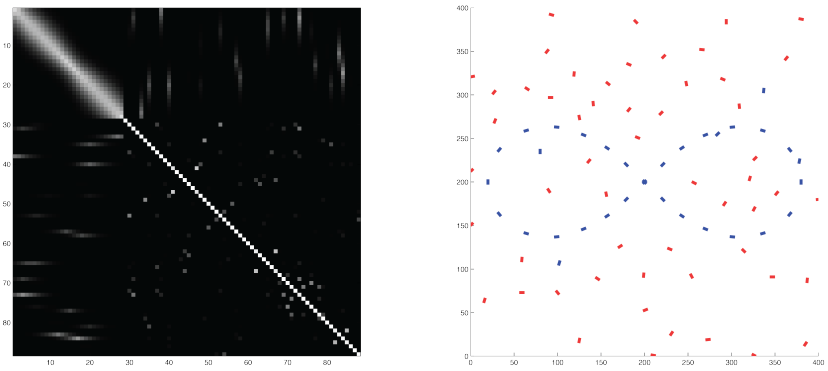

Figure 7: Affinity matrix and grouping results of a lemniscate with random noise.

elements, approximately co-circular with the beginning of the line, are uncorrectly interpeted as being part of the object. Furthermore, the semi-circular contour gets again over-segmented, and many of the randomly collinear points, very far from each other, are interpreted as being a perceptual unit.

In any case, we would like to stress that these kernels do not prevent contours having multiple orientations in the same spatial position to be recognized as one object. See for example Fig. 7, with a lemniscate embedded in a field of randomly oriented segments. On the lifted space where the kernel lives, this 8-figure contour is indeed continuous and non-overlapping, so that the grouping algorithm will assign all of its elements to the same perceptual unit, as human vision would tend to do.

3.2.2 Time asymmetry: a directed graphs approach

In order to set an affinity over a space-time dataset $S_T = \{\eta_i\}_{i=1}^n \subset \mathcal{M}_T$ one must consider the intrinsic notion of causality associated to the time component of η, which from the neural point of view represents the mean activation time of a V1 cell. This causality is reflected in the kernel $\hat{\Gamma}_T^{H,\kappa,\alpha}$, so that any symmetrization would compromise the information content carried by such kernels and negatively affect the efficiency of the grouping algorithm. We will then need to work with a *nonsymmetric affinity matrix* $A_T^{H,\kappa,\alpha}$, hence defining a *directed graph*, given by

$$(a_T^{H,\kappa,\alpha})_{ij} = \hat{\Gamma}_T^{H,\kappa,\alpha}(\eta_i, \eta_j). \tag{17}$$

In order to cope with this asymmetry, we will modify some of the spectral clustering criteria. We will still first normalize the affinity matrix as in (14). This will produce a transition probability matrix P_T, which however is associated to a random walk that does not satisfy the detailed balance condition, because the Markov chain is not anymore reversible. Its eigenvalues $\{\lambda_i\}_{i=1}^n$ and eigenvectors $\{u_i\}_{i=1}^n$ will then be, in general, complex valued. In order to perform spectral clustering in this situation we will follow the approach introduced in [38], which consists concretely in replacing the eigenvectors $\{u_i\}$ with vectors obtained by the sum of their real and imaginary parts $\{u_i^+ = \Re u_i + \Im u_i\}$, and defining their clustering strength in terms of the square modulus of their eigenvalues, hence performing the thresholding step 5. of the algorithm in Table 1 with respect to the set of real numbers $\{|\lambda_i|^2\}_{i=1}^n$.

This is by far not the only possible choice. In particular this choice will not preserve the interpretation of minimal graph cuts, but rather it will produce clusters with the property of having elements with approximately the same outward transition probability. Indeed, as discussed in [34], the minimal cuts clustering and the probabilistic clustering have in general different solutions for directed graphs, which sets a marked difference with the undirected case associated to symmetric affinities.

The choice of working with the couples $\{(u_i^+, |\lambda_i|^2)\}$, that are not in general eigenvalue-eigenvector pairs of a Hermitian matrix, can be better understood as follows. We can see that the eigenvalue problem associated to the real $n \times n$ matrix $P = D^{-1}A$ for a nonsymmetric affinity A can be restated equivalently as an eigenvalue problem with real eigenvectors, by doubling the dimension of the space and replacing P by $\tilde{P} = \begin{pmatrix} P & 0 \\ 0 & P \end{pmatrix}$. In a more formal way, let

$$Pu_\lambda = \lambda u_\lambda \ , \quad P\overline{u_\lambda} = \overline{\lambda}\,\overline{u_\lambda} \tag{18}$$

and call $\lambda = x + iy$. Let us introduce the following notation: set $\vec{u}_\lambda = \begin{pmatrix} u_\lambda \\ \overline{u_\lambda} \end{pmatrix} \in \mathbb{C}^{2n}$, denote with $\tilde{P} = \begin{pmatrix} P & 0 \\ 0 & P \end{pmatrix} = \mathbb{I}_{2\times 2} \otimes P$, and let also $\Lambda_\lambda = \begin{pmatrix} \lambda & 0 \\ 0 & \overline{\lambda} \end{pmatrix} \otimes \mathbb{I}_{n\times n}$, where $\otimes$ is the usual Kronecker product and $\mathbb{I}_{n\times n}$ stands for the $n \times n$ identity matrix. Then solving the problem (18) for $u_\lambda \in \mathbb{C}^n$ is equivalent to solving the problem

$$\tilde{P}\vec{u}_\lambda = \Lambda_\lambda \vec{u}_\lambda \tag{19}$$

for $\vec{u}_\lambda \in \{(z_1, z_2) \in \mathbb{C}^{2n} : z_2 = \overline{z_1}\}$.

Consider now the auxiliary real vectors $\vec{v}_\lambda = \begin{pmatrix} u_\lambda^+ \\ u_\lambda^- \end{pmatrix} = \begin{pmatrix} \Re u_\lambda + \Im u_\lambda \\ \Re u_\lambda - \Im u_\lambda \end{pmatrix}$. If we introduce the matrix $B = \begin{pmatrix} \frac{1-i}{2} & \frac{1+i}{2} \\ \frac{1+i}{2} & \frac{1-i}{2} \end{pmatrix} \otimes \mathbb{I}_{n\times n}$, we have that $\vec{v}_\lambda = B\vec{u}_\lambda$, and we can rewrite (19) as the *real system* for $\vec{v}_\lambda \in \mathbb{R}^{2n}$

$$\tilde{P}\vec{v}_\lambda = \Sigma_\lambda \vec{v}_\lambda \tag{20}$$

where $\Sigma_\lambda = \begin{pmatrix} x & y \\ -y & x \end{pmatrix} \otimes \mathbb{I}_{n\times n} = B\Lambda_\lambda B^{-1} \in \mathbb{R}^{2n\times 2n}$. The magnitude of the action of $\tilde{P}$ over $\vec{v}_\lambda$ is given by $\det \Sigma_\lambda = |\lambda|^2$, which provides then a natural clustering parameter to threshold. Moreover, since the matrix $\tilde{P}$ is a double copy of the matrix P, in order to cluster its rows it is sufficient to work on a one-half dimensional space and choose, without loss of informations, e.g. only the u_λ^+ component of $\vec{v}_\lambda$.

4 Visual grouping with cortical affinities

Several phenomenological findings indicate that the grouping properties obtained by spatial collinearity can easily be broken if one associates a speed and an orthogonal direction of movement to each oriented segment. Limits have been found on the maximum rate of change of local speed along a contour that makes possible the perception of boundaries and shapes. Indeed, a random speed distribution over a dashed line could completely destroy the perception of a single unit as a whole, while enhancing the impression of different segments pertaining to the random background field.

These observations have led to the notion of *motion contour* in [42], where it is shown by psychophysical experiments how the brain groups features together also relying on the local speed perpendicular to their orientation axis, with coherent velocities being represented by velocity fields that vary smoothly over space. The former study expanded the already known notions that local stimulus velocity is discernible (thus determinant for grouping purposes) only when it is orthogonal to the perceived contour or it is not part of a trajectory [20, 30, 55]. Coherently, the analysis carried out in [11] over a data set of cortical neurons in the primary visual cortex showed how the spatio-temporal shape of their RPs is biased to optimally measure the local stimulus orthogonal velocity. Thus, it may be inferred that stimulus local direction of movement and speed are additional features driving the spatial integration involved in the perception of shapes and contours.

The simulations performed in this section will deal with the geometric connectivities constructed in Section 2, which aim to model the neural connections in the cortical area V1 of the visual cortex [3]. These connectivities are defined in feature spaces which include the local orthogonal velocities, which are detected by V1 specialized cells.

4.1 Grouping with spatial features

The perceptual bias towards collinear stimuli has classically been associated to the long-range horizontal connections linking cells in V1 having similar preferences in stimulus orientation. This specialized form of intra-striate connectivity pattern is found across many species, including cats [26], tree shrews [8], and primates [1], the main difference being the specificity and the spatial extent of the connections. Furthermore, axons seem to follow the retinotopic cortical map anisotropically, with the axis of anisotropy being related to the orientation tuning of the originating cell [6]. The clustering algorithm with the affinity matrix $A_3^{H,\kappa}$ constructed as in (16) using the kernel $\hat{\Gamma}_3^{H,\kappa}$ may then be neurally motivated by the assumption that spatial integration, grouping and shape perception are fundamentally modulated by the position and the orientation of the elements in the visual space. Other prominent features of the visual stimulus, namely its velocity and direction of motion, seem to play an important role in the spatial integration of oriented elements [42]. In addition to stimulus orientation, cells in the striate areas are also selective for the direction of motion orthogonal to the cell's preferred orientation [14, 2]. These selectivities are also structurally mapped in the cortical surface, with nearby neurons being tuned for similar motion direction [60], and it has been shown that excitatory horizontal connections in the V1 of the ferret are strictly iso-direction-tuned [45]. In order to model the grouping effect of these connectivities we will then make use of the affinity matrix $A_0^{H,\kappa,\alpha}$ constructed with the extended kernel $\hat{\Gamma}_0^{H,\kappa,\alpha}$.

The stimuli. We have considered a dataset made of n points $\{\zeta_i\}_{i=1}^n$ in the feature space of positions and orientations $\mathcal{M}_3 = \mathbb{R}^2 \times S^1$, having the following structure. Two perceptual units consisting of segments aligned along circle arcs with curvature k are embedded in a background environment of r segment having random positions and orientations. Such dataset will be denoted as S_r^k, and is depicted in Fig. 8, top.

By assigning to each point ζ_i an instantaneous orthogonal velocity v_i, such dataset can be considered as n points $\{\xi_i = (\zeta_i, v_i)\}_{i=1}^n$ embedded in the larger feature space $\mathcal{M}_0 = \mathbb{R}^2 \times S^1 \times \mathbb{R}^+$. A velocity field that is constant on each perceptual unit would describe a rigid motion, and the analysis of such a case with the affinity $A_0^{H,\kappa,\alpha}$ can be understood as an extension of the static case, clearly providing grouping improvement.

A more interesting case is that of a velocity field that changes along each perceptual unit, which represents *shape deformation*. We have then generated a distribution of the velocity feature in the following way: given a fixed maximal velocity V, we have assigned to points belonging to the perceptual units a velocity feature v_i that varies sinusoidally along the arc circle, passing from zero to V, and assigned a random velocity between 0 and V to the background points. Such dataset $\{\xi_i\}_{i=1}^n \subset \mathcal{M}_0$ is depicted in Fig. 8, bottom.

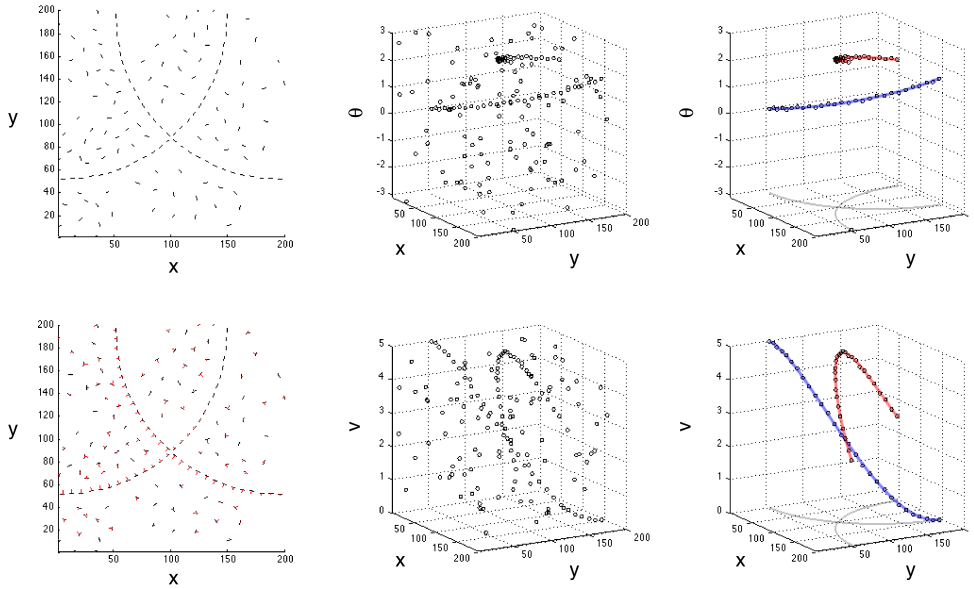

Figure 8: An instance of the dataset S_r^k, where the left column shows the corresponding visual stimulus, the central column shows the dataset in the feature space, and the right column shows the perceptual units, which constitute the target segmentation objects, that for this instance have curvature $k = 0.023$. The top row shows the dataset in the feature space $\mathcal{M}_3 = \mathbb{R}^2 \times S^1$, while the bottom row shows it in the feature space $\mathcal{M}_0 = \mathbb{R}^2 \times S^1 \times \mathbb{R}^+$ with a velocity assignment producing a shape deformation. The bottom left image contains the spatial stimulus (black) with the magnitude of its istantaneous orthogonal velocity (red).

The numerical experiments. We have applied the method described Section 3, using the same spectral clustering parameters used for the examples of Fig. 6, i.e. $\epsilon = 0.05$, $\tau = 150$ and $M = 3$, on a set of stimuli S_r^k for different values of k and r. We have used both the affinity $A_3^{H,\kappa}$ on $\mathcal{M}_0 = \mathbb{R}^2 \times S^1$, and $A_0^{H,\kappa,\alpha}$ on $\mathcal{M}_0 = \mathbb{R}^2 \times S^1 \times \mathbb{R}^+$, in this case setting the maximum velocity assignment to $V = 5$. We have then evaluated the grouping performances for various kernel parameter sets (H, κ) and (H, κ, α) by computing, for each iteration, a percentage error measure

$$E = \frac{E_1 + E_2 + E_3}{n}$$

where n is the total number of points in the stimulus, E_1 is the number of points that were incorrectly assigned to the noise/background set, E_2 is the number of random points that were incorrectly recognized as part of a perceptual unit, and E_3 is the points pertaining to an over- or under-partitioned contour. In order to correctly compare results obtained by partitioning the dataset with a different number of random points,

we have then averaged E over 100 repetitions, where at each repetition we have changed the random part of the stimulus and calculated new kernels $\hat{\Gamma}_3^{H,\kappa}$ and $\hat{\Gamma}_0^{H,\kappa,\alpha}$. The resulting mean percentage error measure $\hat{E}$ has been finally taken as a measure of the quality of the grouping.

The results. The results of the experiments are shown in Fig. 9 and Fig. 10.

In Fig. 9 (a) we have represented with a grayscale intensity the error $\hat{E}$ for the spectral clustering of $S_{120}^{0.056}$ with the affinity $A_3^{H,\kappa}$ in $\mathbb{R}^2 \times S^1$ by varying the parameters (H, κ) of the kernel. The parameters $(k, r) = (0.056, 120)$ correspond to the stimulus in the feature space of positions and orientations composed of the perceptual units with the highest curvature, and having the highest number of random elements, among all the ones tested in the present analysis. A first significative feature emerged is that the set of kernel parameters which gives the lowest grouping error value is $\kappa = 0.056, H = 20$, that is, the same curvature value of the contours in the stimulus and short stochastic path length. Moreover, we observe that by maintaining the kernel parameter H set to its minimum value and decreasing the kernel diffusion coefficient κ we have had an approximately constant higher error value $\hat{E}$, whose dominating components are E_1 and E_3. This indicates that a reduction in the width of the fan of stochastic curves generating Γ_3 impairs the connectivity between high curvature contour elements, so that the algorithm perceives them as separate units and assigns part or all of them to the background/noise set. It is also worth noting that, regardless of the kernel diffusion coefficient, increasing the parameter H negatively impacts the quality of the grouping, this time mainly because of the error component E_2. This was predictable, as longest stochastic path lengths can generate affinities between elements very distant from each other. A high number of random elements can in this case induce the algorithm to recognize them as part of a distant contour. This gives its worse effects when both κ and H have high values, because with the associated kernel a point in $\mathcal{M}_3$ can potentially be connected to distant elements having a very different orientation. In such cases, also the error component E_3 gives its contribute, as the two perceptual units can be under-partitioned and interpreted as one unique object, having each one reciprocally affine contour elements.

Fig. 9 (b) resumes the same kind of grouping analysis carried out with 30 different stimuli $S_r^k \in \mathcal{M}_3$ by varying the parameters k and r. We can notice a positive correlation between the stimulus contour curvature k and the smallest value of the kernel diffusion coefficient κ that provides optimal grouping. Smaller values of κ mostly lead to over-partitioned or unrecognized contours, while higher κ's generally do not impair the contour grouping capability of the algorithm, even if we have observed that the error component E_2 tends to grow together with the number of random elements in the space r and the length of the stochastic paths H. In general thus, in order to correctly identify contours-like objects, the algorithm should be based on a connectivity whose diffusion coefficient is above the contour's curvature.

These results suggest that for a successful grouping the parameters κ and H could be associated to some property of the image detected by the visual cortex. In the discrete case for example, the model could be extended by adding two fiber variables, namely local stimulus curvature and scale, governing the numerical value taken by the two connectivity parameters. One could also argue - and our results are consistent with this view - that curvature and scale are very close concepts, strongly influencing each other, so that κ and H should not be independent. A deeper study of these aspects will be addressed in a future paper.

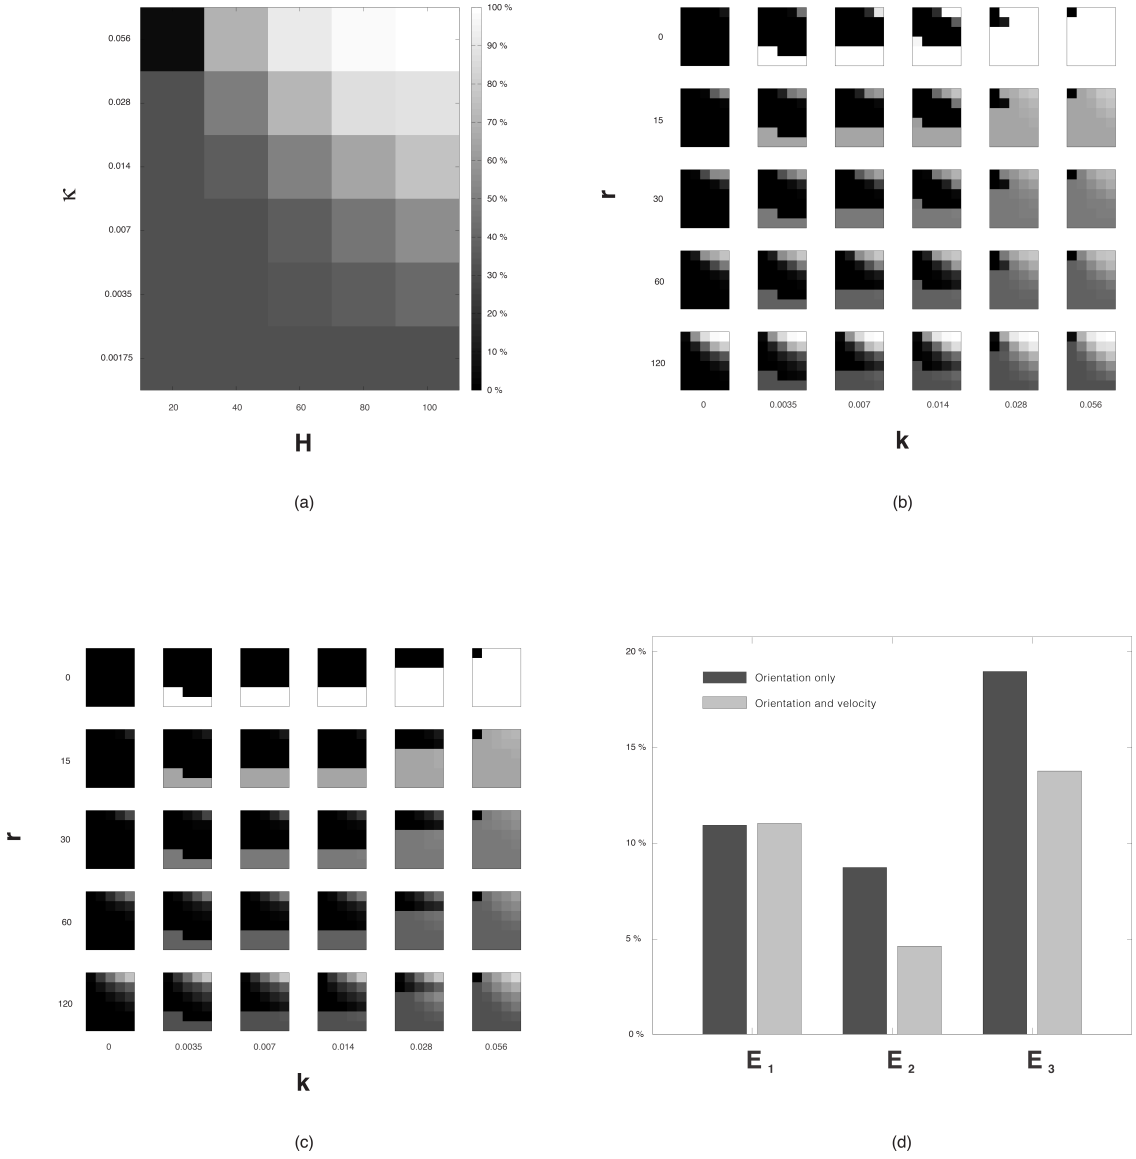

Figure 9: Parametric analysis for visual grouping in $\mathbb{R}^2 \times S^1$ and $\mathbb{R}^2 \times S^1 \times \mathbb{R}^+$. The color intensity of each region is proportional to the percentage of misinterpreted points averaged over 100 repetitions. (a) Grouping results for the stimulus $S_{120}^{0.056}$ in $\mathbb{R}^2 \times S^1$, by varying the kernel parameters (κ, H). (b) Grouping results as in (a) for the stimuli S_r^k in $\mathbb{R}^2 \times S^1$ by varying the stimulus parameters (k, r). (c) Grouping results for the stimuli S_r^k in $\mathbb{R}^2 \times S^1 \times \mathbb{R}^+$ with nonuniform velocity assignments. (d) Comparison of the grouping performances in position-orientation vs position-orientation-velocity with separated analysis of the three error sources.

In Fig. 9 (c) we show the grouping results obtained by analyzing the stimuli S_r^k embedded in $\mathcal{M}_0$, with the described non uniform velocity assignment, using the connectivity $\hat{\Gamma}_0^{H,\kappa,\alpha}$. From a first visual inspection, we can see that the correlation constraint between the contour curvature k and the orientation diffusion coefficient κ is still present. However, the error at the highest values of κ and H is significantly reduced, if not almost completely eliminated for the stimuli having fewer random elements, due to the influence on the algorithm of correlations on the direction of motion. We have not presented the detailed analysis on the parameter α, which is very similar to the one related to κ. We limit ourselves to observe that the best results were obtained for values of α close to $\overline{\alpha} = \dfrac{\pi}{L}V$, where L is the length of the perceptual unit (circle arc) and V is the maximum velocity assigned, which was set to $V = 5$. The reason can be explained in terms of the velocity assignment, which as said we have chosen to be sinusoidally varying along the contour, namely

$$v(\phi) = V \sin\left(\frac{\pi\phi}{L}\right) + \text{const.}$$

where $\phi \in [0, L]$ is the arc parameter of the circular curves. Then $\overline{\alpha} = \max|\dot{v}|$, which is a diffusion constant that ensures that grouping capabilities will be preserved even at the maximum local speed change rate along the contour. In particular, the results shown refer to clustering with the parameter α set to $\overline{\alpha}$.

A comparison of the grouping quality obtained with the orientation connectivity $\hat{\Gamma}_3^{H,\kappa}$ and with the the orientation and velocity connectivity $\hat{\Gamma}_0^{H,\kappa,\alpha}$, with separate presentation of the three error sources considere, is shown in Fig. 9 (d). The histograms correspond to the average of the three error components E_1, E_2 and E_3 over all the stimuli, the kernel parameters and the repetitions. The analysis of E_1 shows that both kernels tend to confuse the perceptual units with noise approximately in the same way. However, from the analysis of E_2 we can see that information on local velocity significantly enhances the performance of the algorithm in the presence of noise, improving the assignment of the random elements to the background. Also the analysis of E_3 shows that the connectivity $\hat{\Gamma}_0^{H,\kappa,\alpha}$ reduces the under-partitioning effect of the stimulus, which typically occurs when the parameters (κ, H) make $\hat{\Gamma}_3^{H,\kappa}$ too long range and widespread.

Finally, some particular cases of the discussed grouping simulations are shown in Fig. 10, where one can see how the use of $\hat{\Gamma}_0^{H,\kappa,\alpha}$ instead of $\hat{\Gamma}_3^{H,\kappa}$ concurs in reducing different kinds of grouping errors of the algorithm.

4.2 Spatio-temporal grouping

In this subsection we consider grouping performed in space and time. Similarly to what happens for the integration of spatial visual information, our visual system is capable to easily predict stimulus trajectories, and to group together elements having similar motion or apparent motion paths. At the psychophysical level, the facilitation in detecting moving stimuli, given a previous cue with a coherent trajectory, is found to be significantly higher than the one expected from the temporal response summation given by the onset and offset dynamics of classical RPs [53]. One possible explanation could be the existence of a specialized facilitatory network linking cells anisotropically and coherently with their direction of motion.

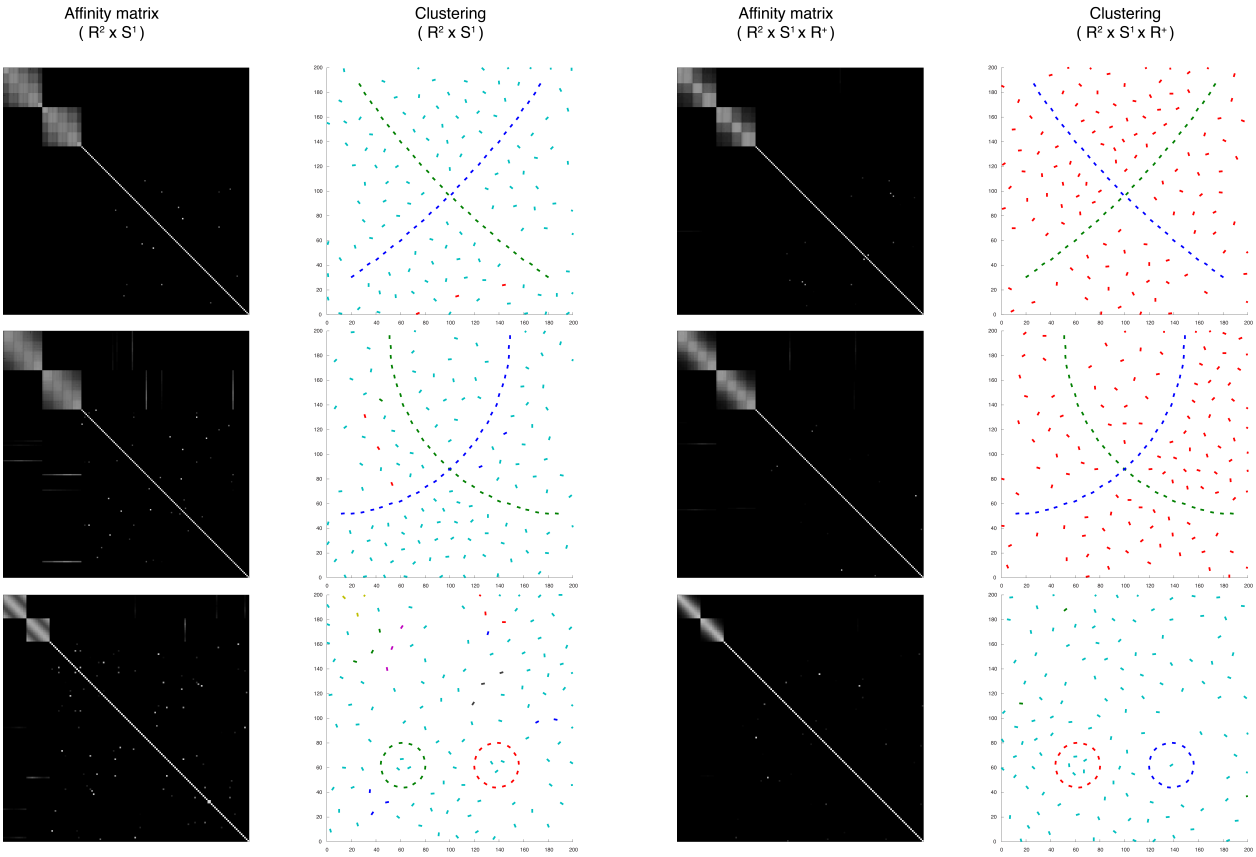

Figure 10: Results of the proposed algorithm for data set with perceptual units with different curvatures. Left column: results obtained by using the position-orientation connectivity given by Γ_3. Right column: results obtained by using the position-orientation-velocity connectivity given by Γ_0. Affinity matrices built with Γ_0 are cleaner than those built with Γ_3: they generally avoid spurious affinities between perceptual units and noise, while maintaining the approximate object block diagonal structure.

In [29], the perception of spatial contours defined by non-oriented stimuli moving coherently and tangentially along a path was studied. The authors suggested rules similar to the ones driving facilitation in position/orientation and inferred a possible role played by a trajectory-specialized network. Another possible evidence of a trajectory-driven connectivity comes from a recent study of the dynamics of neural population response to sudden change of motion direction, where it is shown that for low angular changes a non-linear part of the response provides a sort of spatio-temporal interpolation [63].

We recall also that, at the physiological level, although the basic mechanisms of velocity detection are already present in V1, the estimation of stimulus motion has been classically associated with neurons in visual area MT/V5. There, cells with high selectivity in direction of movement and extended sensitivities to a wide range of stimulus velocities are indeed present [32, 44]. Extra-striate areas are retinotopically organized, with anisotropic and asymmetric connectivity bundles reaching columns of cells tuned for similar orientation and direction preference [31].

On the other hand, the important role played by V1 collinear horizontal connectivity

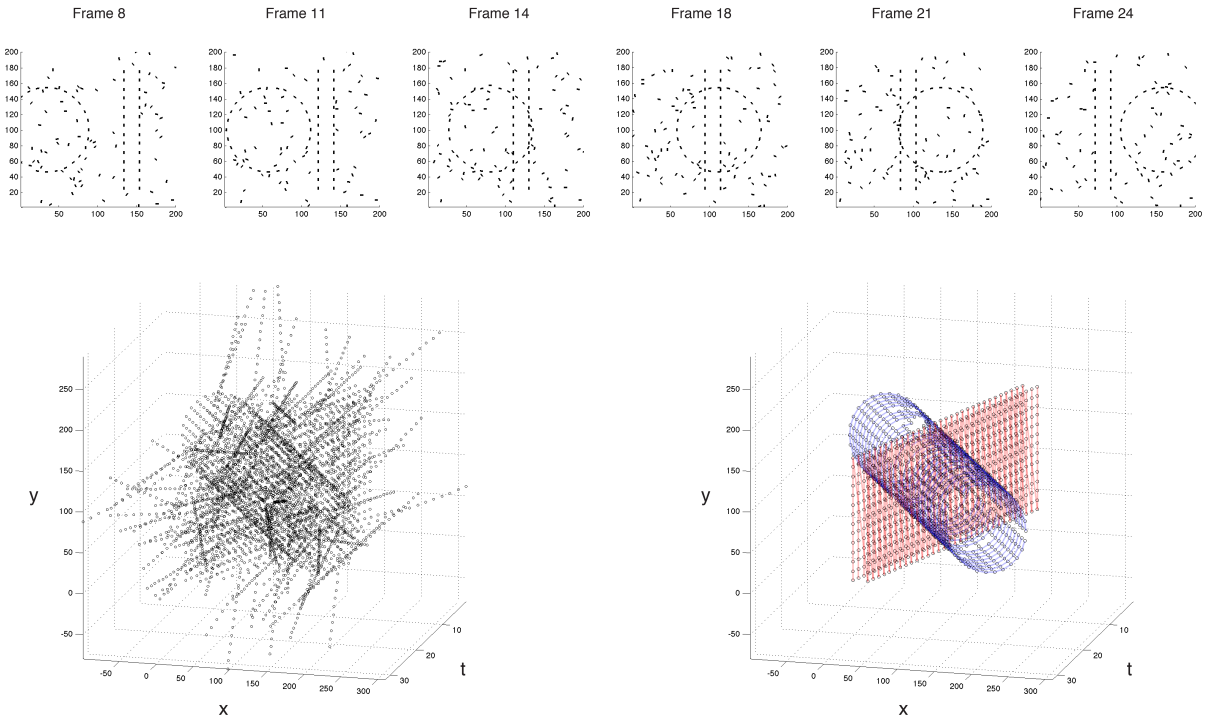

Figure 11: An instantiation of the stimulus living in $\mathbb{R}^2 \times \mathbb{R}^+ \times S^1 \times \mathbb{R}^+$ used in the tests for spatio-temporal grouping of moving shapes. The circle has curvature $k = 0.02$ and its velocity is 7.5 units/frame, that is twice that of the bars (3.75 units/frame).

in motion perception has been suggested already in [50], after showing that perception of speed is biased by the direction of motion. Also, recent results in [37] showed how trajectories of oriented segments are significantly more detectable for orientations orthogonal to the path of motion, thus supporting the hypothesis of two different facilitatory mechanisms [20].

The stimulus. The dataset that we have used in this simulation is depicted in Fig. 11. We created a set of points $\eta_i = (x_i, y_i, t_i, \theta_i, v_i) \in \mathcal{M}_T = \mathbb{R}^2 \times \mathbb{R}^+ \times S^1 \times \mathbb{R}^+$, represented in the figures by segments moving in time for a total of $n_t = 32$ frames. Three perceptual units, a circular shape of curvature $k = 0.02$ and two bars, move as rigid bodies, with the circle translating with speed $V_{\text{circ}} = 7.5$ spatial units/frame in the opposite direction to the bars, who both move with speed $V_{\text{bars}} = 3.75$ spatial units/frame. We recall that the fiber coordinate of local velocities v_i of each point represents just the projection onto $\vec{X}_3$, i.e. orthogonal to the segment orientation, of the real velocity vector. Similarly to what we did in the previous example, these perceptual units were embedded into a background consisting of a variable number r of random elements, each one having a uniform motion path along their $\vec{X}_3$ direction during all the stimulus frames.

The numerical experiment. The aim of the grouping algorithm that we will use for this experiment is to carry out a segmentation of the full spatio-temporal surfaces representing the moving objects.

From the point of view of the structure of the stimulus which gives rise to the dataset, this represents a composite task. It consists of a visual grouping at the spatial level, which identifies the objects, and a visual grouping at the spatio-temporal level, where each previously identified cluster is recognized as constituting the same object during its movement. It is then reasonable to assume that the two connectivity mechanisms $\hat{\Gamma}_0^{H,\kappa,\alpha}$, modeling the interactions between points of a motion contour, and $\hat{\Gamma}_T^{H,\kappa,\alpha}$, which models motion integration of point trajectories, combine in the clustering of spatio-temporal perceptual units.

On the other hand, the presence of more than one grouping law governing the detection of contours, with different underlying implementing structures, was experimentally confirmed at the psychophysical level in [30]. Moreover, such composition of different mechanisms resulted to be compatibile with a probabilistic summation.

Guided by such arguments, we have then chosen to perform the spectral clustering on this dataset with a matrix P obtained as the sum of the transition probability (normalized affinity) matrices obtained from the cortical connectivities $\hat{\Gamma}_0^{H,\kappa,\alpha}$ and $\hat{\Gamma}_T^{H,\kappa,\alpha}$. More precisely, we have constructed the symmetric matrix A_0, based on $\hat{\Gamma}_0^{H,\kappa,\alpha}$ as in (16), with the additional condition of setting zero affinity between points having different temporal coordinates, and we have constructed the nonsymmetric matrix A_T, based on $\hat{\Gamma}_0^{H,\kappa,\alpha}$ as in (17). We have then normalized both of them as in (14), obtaining the transition probabilities P_0 and P_T, and we have defined the combined spatio-temporal normalized affinity as

$$P = \frac{P_0 + P_T}{2}. \tag{21}$$

From the neural point of view, we observe that by relying on the normalized affinity (21) we are implcitly assuming a much faster propagation along the connectivity defined by $\hat{\Gamma}_0^{H,\kappa,\alpha}$, with respect to the temporal dynamics. More precisely, by assigning zero A_0 affinity to temporally separated points of the dataset, we are considering spatial connections that fully operate at each single frame, which corresponds to an almost instantaneous velocity detection and a high horizontal transmission speed.

Finally, in order to perform the spectral clustering over P, we will use the approach described in Section 3.2.2. Indeed, since A_T is not symmetric, P will have in general complex eigenvalues and eigenvectors, but it was constructed in such a way to keep a probabilistic structure.

The results. The results obtained for various instances of the circle/bars stimulus by varying the number of background/noise elements r are shown in Fig. 12.

The parameters chosen to run the algorithm were $\epsilon = 0.01$, $\tau = 150$ and $M = 3$, where a smaller threshold ϵ with respect to the previous experiments is suggested by the presence of a semantically and geometrically sharper affinity. The integration parameter was set to $H = 40$ for both Γ_0 and Γ_T, according to the quality indications of the previous experiment. Similarly, we have set the angular diffusion parameters to the same value $\kappa = 0.014$ for both kernels.

With respect to the diffusion coefficient over velocities, we are showing the results for values that for both kernels stay close to the optimal value for the circle discussed in the previous subsection, that is $\overline{\alpha} = \frac{kV_{\text{circ}}}{2} = 0.75$. Namely, we have set $\alpha = 0.5$ for Γ_0 and $\alpha = 1$ for Γ_T. In contrast to all other parameters, which were chosen to be equal for both kernels, this one is indeed observed to perform better when it is larger on the temporal connectivity. Such a behavior is coherent with the consider-

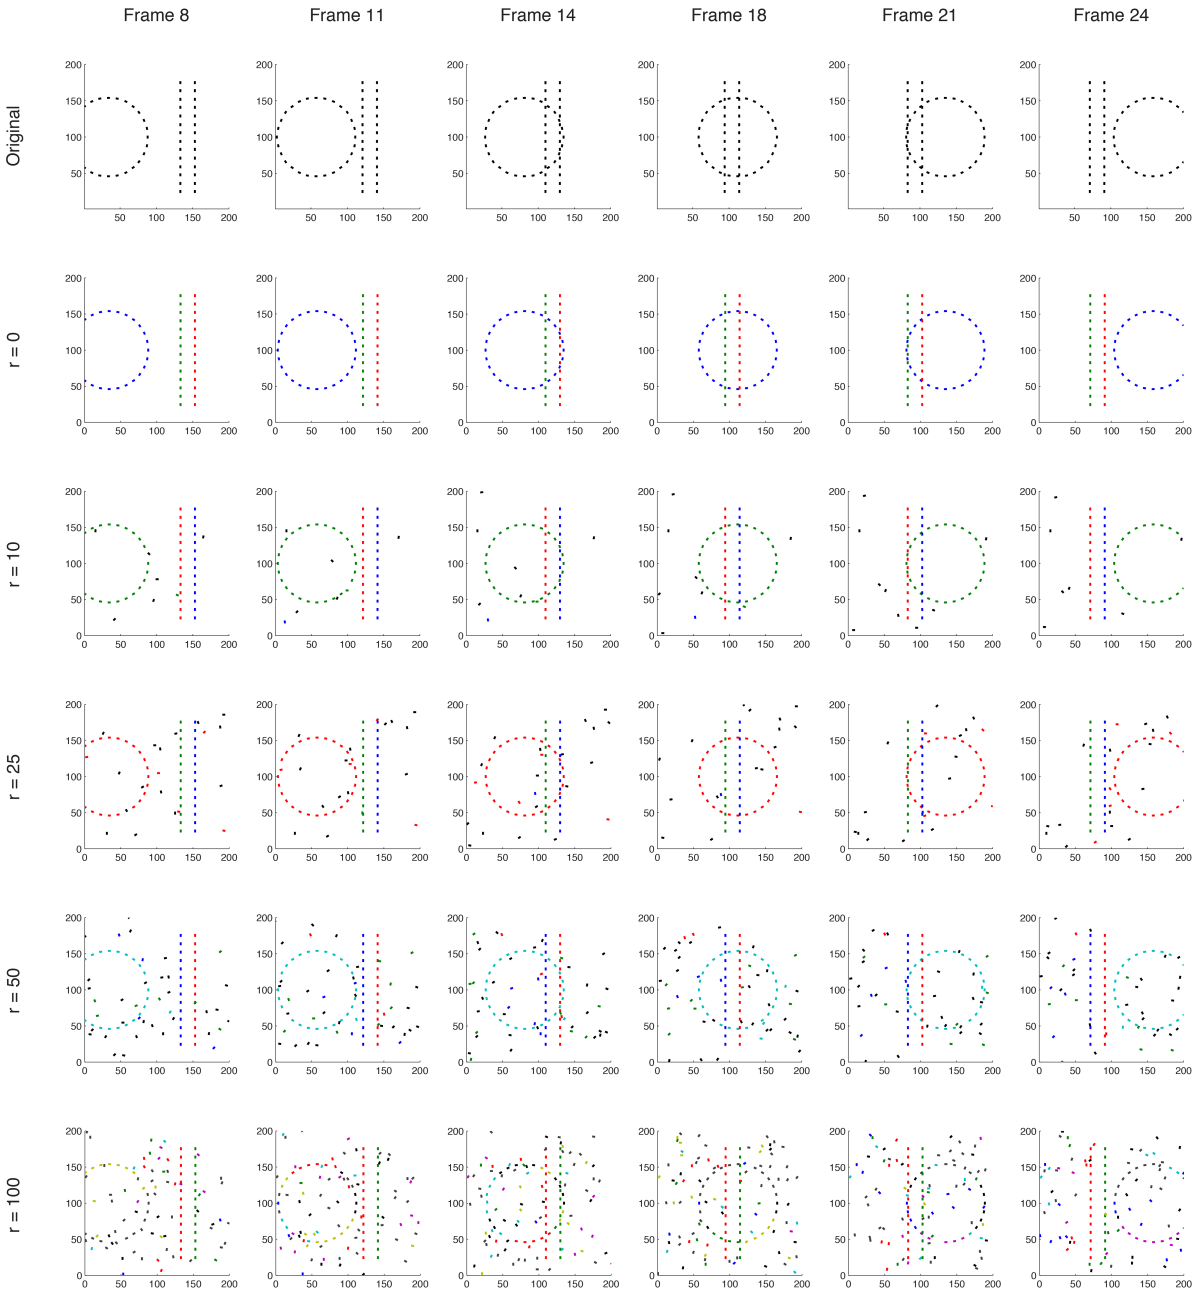

Figure 12: Results obtained by using both connectivities $\hat{\Gamma}_0^{H,\kappa,\alpha}$ and $\hat{\Gamma}_T^{H,\kappa,\alpha}$ simultaneously. The grouping is succesful for stimuli with noise levels up to 50%. At higher noise level the algorithm fails by over-partitioning the countours of the moving circle. The third row shows grouping results with $r = 50$ background elements, corresponding to about the 47% of the total. The fourth row shows grouping results with $r = 100$ background elements, corresponding to about the 64% of the total.

ations made in [20, 55], where the effect of local changes in the velocity of a motion contour over the perception of visual units are studied. The results of these works show that such changes tend to weaken the visual grouping, but this effect is much stronger when changes are orthogonal to contours with respect to changes that are tangential to trajectories. In the present setting, we are performing visual grouping both at the level of motion contours, with Γ_0, and at the level of trajectories, with Γ_T, describing two copresent connectivities, and the parameter α describes the sensitivity of the corresponding kernel to local changes in the velocity, so the different levels of α reflect a higher sensitivity (lower diffusion) to velocity for the Γ_0 connectivity, and a lower sensitivity (higher diffusion) for the Γ_T connectivity.

From Fig. 12 it is possible to see how spectral clustering with the composite affinity (21) performs in recognizing the spatio-temporal surfaces relative to the moving circle and the bars with different noise levels. It is succesful in clustering the perceptual units, separating clearly their boundaries from the background and between themselves, up to relativeli high noise levels. When the number r of random elements was lower than about the 50% of the total, we obtain correct clustering, but for higher noise values the algorithm began to give poor grouping results, as in the shown case of $r = 100$ random segments, which correspond to about 64% of the total. It is worth noting that even at the higher noise values, the bars are always correctly retrieved. The algorithm tends indeed to fail in the detection of contours with high curvature, confusing them with the background segments, thus leading to over-partitioning.

In general, though, we show that the connectivity kernels defined by the proposed cortical-inspired geometrical model applied to a simple spectral clustering algorithm are able to carry out a non-trivial grouping task. To better understand the powerful mechanics involved in the calculations, let's consider for example that the only segments of the circle that present a positive affinity value with their corresponding points at future temporal positions are the ones having an orientation value near $\pm\frac{\pi}{2}$, as only for them the vector field $\vec{X}_5$ of connectivity propagation has the same direction of the global movement of the shape.

In fact, while the ability and the reliability of visual neurons in areas V1 and MT/V5 in measuring local stimulus orientation and speed have been studied extensively, the majority of cells in those areas respond solely to the local characteristics they are tuned for. In doing so, the measurements available in the first stages of the visual cortex are subject to the well-known aperture problem: with no information other than the local direction of movement, it cannot be said much about the real direction and speed of the object to which that local measurement refers. For a continuously moving contour, for example, classical orientation- and direction-selective cortical cells measure, for each position along the contour, just the velocity component that is orthogonal to the contour tangent direction at that point. In the framework presented throughout the paper, this is modeled so that the fiber variable of local velocity refers to movements in the $\vec{X}_3$ direction.

5 Conclusions

In this work we have constructed a clustering algorithm for visual grouping of spatio-temporal stimuli in terms of geometric connectivity kernels associated to the functional architecture of the visual cortex. The main purpose was to test the segmentation capabilities of such a geometric model of low-level vision areas, with respect to spectral

analysis mechanisms. Previous experimental investigations such as [30] already indicated that spatio-temporal perceptual associations play an important role in visual recognition tasks. The recent results in [48] also suggest the presence of concrete cortical implementations of a spectral analysis associated to lateral connection.

We have used recent dimensionality reduction methods [25], chosen for their robustness and for their relatively simple structure, which allowed us to focus primarily on kernel properties. We have then performed several spectral clusterings on the introduced spatio-temporal cortical feature space of position, time, orientation and velocity, by using anisotropic affinities obtained by the models of cortical connections. Such affinities present a structure that is geometrically adapted to the considered stimuli, and proved to be able to better extract the relevant information compared to more classical kernels such as the isotropic Gaussian one. In particular, the proposed algorithm is capable to group together elements belonging to a single contour or shape moving in time, forming a spatio-temporal surface, and to distinguish them from a noisy background.

The first analysis that we carried out considered visual grouping when the affinity between points of a visual stimulus is assigned in the cortical feature space of positions and local orientations $\mathbb{R}^2 \times S^1$, and in the feature space $\mathbb{R}^2 \times S^1 \times \mathbb{R}^+$ taking locally detected velocities into account. The produced algorithm showed generally higher segmentation capabilities with respect to isotropic affinities defined only on the visual stimulus. Moreover, performance analyses showed that when using the additional velocity feature the results are less affected by the influence of random elements and proved significantly lower percentage of grouping errors due to over- or under-partitioning, thus allowing the horizontal connectivity to be spatially extended without suffering noise, as it happens in the visual cortex [6].

A second analysis considered grouping in space and time, and produced an algorithm that is able to identify spatial perceptual units and to follow them during their motion. This is done by extending the previous approach, combining the affinity over $\mathbb{R}^2 \times S^1 \times \mathbb{R}^+$, here treated as an instantaneous connectivity, with an affinity in the cortical feature space of positions, activation times, local orientations and locally detected velocities $\mathbb{R}^2 \times \mathbb{R}^+ \times S^1 \times \mathbb{R}^+$. In order to cope with the intrinsic causality of time evolution, we have worked direcly on asymmetric affinities and used probabilistic clustering arguments [38]. The copresence of more than one connectivity mechanism in the segmentation of spatio-temporal stimuli is actually a realistic assumption on visual cortex behavior [20].

We would like to remark that modelling the neural dynamical aspects of visual perception is a very delicate task. Indeed, many open questions remain to be addressed, regarding in particular how mean-field equations in space-time can be compatible with the fast time scales of visual processing, or how the delays introduced by the neural dynamics influence the functionalities of cortical architectures. This first model of functional architecture in space-time does not take into account the integration time constants for neurons, as well as other biophysical phenemona, which certainly matter and deserve future study. These aspects are also strictly related to the general problem of understanding the time scale of mean-field dynamics at the physical level. We observe however that this model of connectivity is able to implement a preactivation mechanism induced by spatio-temporal stimuli, whose biological plausibility was discussed and compared with neurophisiological measurement in [3], showing good accordance.

The present investigation has thus provided a geometric framework to perform

clustering in feature spaces, following principles of extensions that are feasible to be further generalized to higher dimensional detected features. Possible extensions could be achieved by the inclusion of features such as color, three dimensional stereo, or scale as additionaly detected features. Depending on the geometry of the stimuli, and on the psychophysiological indications, such extensions can lead to the definition of connectivities of higher complexity, or to the addition of copresent association mechanisms, following the two modalities that were discussed. This approach provides both a way to desgin artificial perceptual algorithms adapted to the geometry of the information, and a tool to propose new models of actual connections in the visual cortex, also suggesting further psychophysical or physiological experiments to compare and tune the model's parameters in order to fit real visual perception and cognition behaviors. Moreover, a natural future step will be the inclusion of realistic feature detection mechanisms allowing to apply this methodology to real stimuli.

References

[1] Angelucci, A., Levitt, J.B., Walton, E.J.S., Hupe, J., Bullier, J., Lund, J.S.: Circuits for local and global signal integration in primary visual cortex. J. Neurosci. 22(19), 8633-8646 (2002)

[2] Barbieri, D., Citti, G., Sarti, A.: How Uncertainty bounds the shape index of simple cells. The Journal of Mathematical Neuroscience 4 (5), Special Issue "Uncertainty and the brain" (2014).

[3] Barbieri, D., Citti, G., Cocci, G., Sarti, A.: A cortical-inspired geometry for contour perception and motion integration. J. Math. Imag. Vis. Jan (2014).

[4] Barnett, L., Buckley, C. L.: Neural complexity and structural connectivity. Phys. Rev. E 79. 051914 (2009)

[5] Belkin, M., Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation. Neur. Comp. 15(6), 1373-1396 (2003)

[6] Bosking, W.H., Zhang, Y., Schofield, B., Fitzpatrick, D.: Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. J. Neurosci. 17(6), 2112-2127 (1997)

[7] Bressloff, P.C., Cowan, J.D., Golubitsky, M., Thomas, P.J., Wiener, M.C.: What Geometric Visual Hallucinations Tell Us about the Visual Cortex, Neural Computation, 14 (2002), pp. 473-491.

[8] Chisum, H.J., Mooser, F., Fitzpatrick, D.: Emergent properties of layer 2/3 neurons reflect the collinear arrangement of horizontal connections in tree shrew visual cortex. J. Neurosci. 23(7), 2947-2960 (2003)

[9] Citti, G., Sarti, A.: A cortical based model of perceptual completion in the roto-translation space. J. Math. Imag. Vis. 24(3), 307-326 (2006)

[10] Citti, G., Sarti, A. (Ed): Neuromathematics of Vision. Lecture Notes in Morphogenesis, Springer 2014.

[11] Cocci, G., Barbieri, D., Sarti, A.: Spatio-temporal receptive fields of cells in V1 are optimally shaped for stimulus velocity estimation. J. Opt. Soc. Am. A 29(1), 130-138 (2012)

[12] Coifman, R.R., Lafon, S.: Diffusion maps. Appl. Comp. Harm. Anal. 21, 530 (2006)

[13] Daugman, J., Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. J. Opt. Soc. Am. A 2(7), 1160-1169 (1985)

[14] DeAngelis, G.C., Ohzawa, I., Freeman, R.D.: Receptive-field dynamics in the central visual pathways. Trends Neurosc. 18, 451-458 (1995)

[15] Duits, R., van Almsick, M.: The explicit solutions of linear left-invariant second order stochastic evolution equations on the 2D Euclidean motion group. Quart. Appl. Math. 66, 27-67 (2008)

[16] Faugeras, O., Veltz, R., Grimbert, F.: Persistent neural states: stationary localized activity patterns in nonlinear continuous n-population, q-dimensional neural networks. Neural Comput. 21, 147187 (2009).

[17] Field, D.J., Hayes, A., Hess, R.F.: Contour integration by the human visual system: evidence for a local association field. Vision Res. 33, 173-193 (1993)

[18] Geisler, W.S., Perry, J.S., Super, B.J., Gallogly, D.P.: Edge co-occurrence in natural images predicts contour grouping performance. Vision Res. 41, 711-724 (2001)

[19] Graham, C., Talay, D.: Stochastic siulation and Monte Carlo methods. Springer 2013.

[20] Hess, R.F., Ledgeway, T.: The detection of direction-defined and speed-defined spatial contours: One mechanism or two? Vision Research 43, 597606 (2003)

[21] Hoffman, W.C.: The visual cortex is a contact bundle. Appl. Math. Comput. 32, 137-167 (1989)

[22] Hubel, D.H.: Eye, Brain and Vision. Scientific American Library, New York (1988)

[23] Hubel, D.H., Wiesel, T.N.: Ferrier lecture: Functional architecture of macaque monkey visual cortex. Proceedings of the Royal Society of London Series B 198(1130), 1-59 (1977)

[24] Jones, J.P., Palmer, L.A.: An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. J. Neurophysiol. 58(6), 1233-1258 (1987)

[25] Kannan, R., Vempala, S., Vetta, A.: On clusterings: Good, bad and spectral. Journal of the ACM 51(3), 497-515 (2004)

[26] Kisvarday, Z.F., Toth, E., Rausch, M., Eysel, U.T.: Orientation-specific relationship between populations of excitatory and inhibitory lateral connections in the visual cortex of the cat. Cereb. Cortex 7(7), 605-618 (1997)

[27] Koenderink, J.J., van Doorn, A.J: Representation of local geometry in the visual system. Biol. Cybernet. 55, 367-375 (1987)

[28] Lafon, S., Lee, A. B.: Diffusion maps and coarse-graining: A unified framework for dimensionality reduction, graph partitioning, and data set parameterization. IEEE Trans. Pattern Anal. Mach. Intell. 28(9), 1393-1403 (2006)

[29] Ledgeway, T., Hess, R.F.: Rules for combining the outputs of local motion detectors to define simple contours. Vision Research 42, 653659 (2002)

[30] Ledgeway, T., Hess, R.F., Geisler, W.S.: Grouping local orientation and direction signals to extract spatial contours: empirical tests of "association field" models of contour integration. Vision Res. 45(19), 2511-2522 (2005)

[31] Malach, R., Schirman, T.D., Harel, M., Tootell, R.B., Malonek, D.: Organization of intrinsic connections in owl monkey area MT. Cereb. Cortex 7(4), 386-393 (1997)

[32] Maunsell, J.H.R., Van Essen, D.C.: Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. J. Neurophysiol. 49(5), 1127-1147 (1983)

[33] Meila, M., Shi, J.: A random walks view of spectral segmentation. 8th International Workshop on Articial Intelligence and Statistics (2001)

[34] Meila, M., Pentney, W.: Clustering by weighted cuts in directed graphs. Proceedings of the 2007 SIAM International Conference on Data Mining (2007)

[35] Mumford, D.: Elastica and computer vision. In: Bajaj, C.L. (eds.) Algebraic Geometry and its Applications, pp. 491-506. Springer-Verlag (1994)

[36] Ng, A.Y., Jordan, M.I., Weiss, Y.: On spectral clustering: analysis and an algorithm. Cambridge, MA. Advances in Neural Information Processing Systems 14. MIT Press (2002)

[37] Pavan, A., Casco, C., Mather, G., Bellacosa, R. M., Cuturi, L. F., Campana, G.: The effect of spatial orientation on detecting motion trajectories in noise. Vis. Res. 51, 2077-2084 (2011)

[38] Pentney, W., Meila, M.: Proceedings of the 20th national conference on Artificial intelligence (AAAI) 2, 845-850 (2005)

[39] Perona, P., Freeman, W.T.: A factorization approach to grouping. Proc. ECCV, 655-670 (1998)

[40] Petitot, J., Tondut, Y.: Vers une neurogéométrie. Fibrations corticales, structures de contact et contours subjectifs modaux. Math. Sci. Hum. 145, 5-101 (1999)

[41] Platen, E.: An introduction to numerical methods for stochastic differential equations. Acta Numer. 8, 197-246 (1999)

[42] Rainville, S.J.M., Wilson, H.R.: Global shape coding for motion-defined radial-frequency contours. Vision Res. 45(25-26), 3189-3201 (2005)

[43] Ringach, D., Shapley, R.: Reverse correlation in neurophysiology. Cognitive Science 28(2), 147-166 (2004)

[44] Rodman, H.R., Albright, T.D.: Coding of visual stimulus velocity in area MT of the macaque. Vis. Res. 27(12), 2035-2048 (1987)

[45] Roerig, B., Kao, J.P.: Organization of intracortical circuits in relation to direction preference maps in ferret visual cortex. J Neurosci. 19(24) (1999)

[46] Roweis, S., Saul, L.: Nonlinear dimensionality reduction by locally linear embedding. Science, 290(5500), 2323-2326 (2000)

[47] Sanguinetti, G., Citti, G., Sarti, A.: A model of natural image edge co-occurrence in the rototranslation group. J. Vis. 10(14) (2010)

[48] Sarti, A., Citti, G.: The constitution of visual perceptual units in the functional architecture of V1. Submitted. arxiv.org/abs/1406.0289

[49] Sato, K.-I.: Lévy processes and infinitely divisible distributions. Cambridge University Press 1999.

[50] Series, P., Georges, S., Lorenceau, J., Fregnac, Y.: Orientation dependent modulation of apparent speed: a model based on the dynamics of feed-forward and horizontal connectivity in V1 cortex. Vis. Res. 42(25), 2781-2797 (2002)

[51] Shi, J., Malik, J.: Normalzied cuts and image segmentation. Proc. IEEE Conf. Comp. Vis. and Pattern Recogn, 731-737 (1997)

[52] Tononi, G., Sporns, O., Edelman, G. M.: A measure for brain complexity. Relating functional segregation and integration in the nervous system. Proc. Natl. Acad. Sci. USA 91, 5033-5037 (1994)

[53] Verghese, P., Watamaniuk, S.N.J., McKee, S.P., Grzywacz, N.M.: Local motion detectors cannot account for the detectability of an extended trajectory in noise. Vision Res. 39(1), 19-30 (1999)

[54] Verghese, P., McKee, S.P., Grzywacz, N.M.: Stimulus configuration determines the detectability of motion signals in noise. J. Opt. Soc. Am. A 17(9), 1525-1534 (2000)

[55] Verghese, P., McKee, S.P.: Motion grouping impairs speed discrimination. Vis. Res. 46(8-9), 1540-1546 (2006)

[56] Von Luxburg, U.: A tutorial on spectral clustering. Statistics and Computing 17(4), 395-416 (2007)

[57] Wagemans, J., Elder, H. J., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., von der Heydt, R.: A Century of Gestalt Psychology in Visual Perception: I. Perceptual grouping and Figure-Ground Organization. Psychol Bull. 138(6), 1172-1217 (2012)

[58] Watamaniuk, S.N.J.: The predictive power of trajectory motion. Vision Res. 45(24), 2993-3003 (2005)

[59] Weiss, Y.: Segmentation using eigenvectors: a unifying view. Proc. IEEE Int. Conf. on Comp. Vis, 975-982 (1999)

[60] Weliky, M., Bosking, W.H., Fitzpatrck, D.: A systematic map of direction preference in primary visual cortex. Nature 379(6567), 725-728 (1996)

[61] Wertheimer, M.: Laws of organization in perceptual forms. Ellis, Willis D. (Ed). A source book of Gestalt psychology, 71-88 (1938)

[62] Williams, L.R., Jacobs, D.W.: Stochastic completion fields. ICCV Proceedings (1995)

[63] Wu, W., Tiesinga, P.H., Tucker, T.R., Mitroff, S.R., Fitzpatrick, D.: Dynamics of population response to changes of motion direction in primary visual cortex. J. Neurosci. 31(36), 12767-12777 (2011)

[64] Zelnik-manor, L., Perona, P.: Self-tuning spectral clustering. Advances in Neural Information Processing Systems 17, 1601-1608 (2004)

[65] Zhang, J., Duits, R., ter Haar Romeny, B. M.: Numerical Approaches for Linear Left-invariant Diffusions on SE(2), their Comparison to Exact Solutions, and their Applications in Retinal Imaging. Preprint. arxiv.org/abs/1403.3320

Giacomo Cocci
DEI - Department of Electrical, Electronic, and Information Engineering "Guglielmo Marconi", University of Bologna. Viale del Risorgimento 2, 40136 Bologna, Italy.
E-mail: giacomo.cocci2@unibo.it
Davide Barbieri
Department of Mathematics, Autonomous University of Madrid. Facultad de Ciencias, Campus Cantoblanco, 28049 Madrid, Spain.
E-mail: davide.barbieri@uam.es
Giovanna Citti
Department of Mathematics, University of Bologna. Piazza di Porta San Donato 5, 40126 Bologna, Italy.
E-mail: giovanna.citti@unibo.it
Alessandro Sarti
CAMS - Centre d'Analyse et de Mathématique Sociales. EHESS, 190-198 Avenue de France, 75244 Paris, France.
E-mail: alessandro.sarti@ehess.fr